



Intervening to Learn About Effects of One's Actions

Robert Ian Bowers^{1,2*} and William Timberlake¹

¹ Indiana University

² School of Psychology, University of Minho

*Corresponding author (ribowers@gmail.com)

Citation – Bowers, R. I., & Timberlake, W. (2018). Intervening to learn about effects of one's actions. *Animal Behavior and Cognition*, 5(1), 79-102. <https://doi.org/10.26451/abc.05.01.07.2018>

Abstract – Intervention provides a reliable cue to veridical causality. Just as an experimenter manipulates variables to distinguish cause and effect from mere correlation, so might a rat learn differently about the effects of its own actions. However, theory remains vague on specific predictions. The present study asks whether and how producing a conditional stimulus by lever contact alters what a rat learns about that stimulus. Despite the theoretical pressure among theories of causal reasoning for an effect of intervention to hold, the effect we found was not in standard goal-oriented response variables, but in a general activity measure, and so not readily interpretable by typical theories of causal reasoning. We propose a viable explanation for this pattern in terms of multiple foraging strategies. The primary contribution of the present results is that they pose three challenges to theories that concern how animals deal with cause and effect: 1) to resolve present ambiguities regarding predictions; 2) to situate causal cognition research in specific ecological contexts, such as predation; and 3) to look beyond goal acquisition, to the rest of the animal's behavior, including general activity. We present one potential solution in alignment with behavior systems approaches.

Keywords – Intervention, Causal reasoning, Graphical models, Behavior systems, General activity, *Rattus norvegicus*

Experimenters manipulate variables to discern cause and effect from mere correlation. They understand that producing a potential cause by one's own action to observe its effects (intervention) is crucial to their ability to attribute causation, and how observing covariation in this context provides much stronger evidence of causation than observing the same pattern occurring without control over relevant variables. However, it is not only experimenters who experiment. Cooks experiment with recipes; electricians experiment with circuits; and the rest of us experiment with recipes, circuits, plumbing, automotive repair, and whenever we attempt to decipher hidden causal factors. When we do, we are not always very principled scientists about it (even the principled scientists), but whatever experimental control we may overlook while adjusting the amount of sugar in a cake, reliably, when we test causal factors, here or in any domain, we intervene on variables.

Likewise, rats are not merely passive observers, but engage in reliable patterns of appetitive behavior in which they interact with objects in their environments in system-specific ways. A feeding-ready rat will root, gnaw, and manipulate protrusions, thereby producing effects that may provide meaningful cues to causality. Such information is often available to be exploited by an active learner. Do they use this information appropriately? Are rats naïve experimenters? To be clear, the question is not whether rats are doing what a trained experimenter does. Rather, one can imagine a rat testing hypotheses with no more foresight than shown by human naïve experimenters. For instance, consider a man catching a glimpse of the glimmer of his watch reflecting the sun's light onto a wall. He would not be sure of the

cause until he intervened, manipulating the angle of his watch and finding the movement of the flash of light to follow, perhaps just once or twice. To do this requires no knowledge of scientific method. Intervention works for the naïve experimenter just as it does for the trained experimentalist, and for the same reasons. Intervention provides the active learner a reliable cue to veridical causality.

By some views, the capacity of intervention to elucidate causation presupposes that the actor entertains the hypothesis tested. The analogy with experimentation encourages this assumption, since holding a hypothesis to test is an important part of the experimental method. Thus, benefiting from intervention may appear to imply possession of beliefs about potential causal factors, and thereby appear to oblige a metaphysically weighty conception of intervention. However, the entity filling the role of the experimenter's hypothesis need not be representational. In principle, prior learning, or the structure of motivational systems could provide such bases.

How does intervention affect learning about manipulated stimuli? If rats use intervention as a source of information about the predictive value of a stimulus, those with the opportunity to intervene may show signs of enhanced learning. In the present experiments, pairs of subjects in yoked chambers each periodically receive a stimulus followed by food delivery. This develops and maintains responding to the stimulus. In addition, one subject's (INT) lever contact produces the same stimulus at any moment in both chambers. Her partner (OBS) also has a lever but cannot produce that stimulus herself. Thus, both subjects receive the same pattern of conditional stimulus presentation. The difference is that one subject produces the stimulus by her own actions. Of interest is whether and how patterns of conditional responding to these two stimuli differ, and what this indicates about the effects of intervention on learning. When intervening subjects press, they and their partners are learning something about the stimulus, but is the learning of these two animals the same?

Note that our question regarding a role of intervention in learning is orthogonal to questions concerning learning about the outcomes of instrumental responses, to which much research has been directed (e.g., Thorndike, 1911). Other studies with rats pressing levers have been phrased in terms of intervention, but there is variety to how intervention has been conceptualized, and to its assumed significance to learning. For instance, Polack, McConnell, and Miller (2013) "assumed that changes in operant responding consistent with the subject's motivation reflected goal-directed behavior, which [they] equated with intervention, and therefore constituted behavior indicative of causal knowledge" (p. 27-28). This approach has had a very influential precedent in the thinking of Shanks and Dickinson (e.g., 1991; Dickinson, 2001). Polack et al. (2013) showed that instrumental responding tracks outcome devaluation (e.g., Adams & Dickinson, 1981) with an aversive stimulus. However, granting their assumption circumvents the questions about intervention the present study attempts to address. Blaisdell, Sawa, Leising, and Waldmann (2006) similarly spoke of lever press as intervention, but that study conceptualized lever press no differently than other hypothetical causes. Leising, Wong, Waldmann, and Blaisdell (2008) stressed that they were showing a role for intervention, but based that claim on the dubious assumption that rats would perceive a novel event as a cause of a following stimulus even without prior training, and presented this in contrast to a replication of Blaisdell et al.'s (2006) demonstration (also readily attributable to differences in stimulus salience; see Burgess, Dwyer, & Honey, 2012, for a critical reply). In contrast to these prior studies, the present question concerns learning about a conditional stimulus that a rat exerts control over by pressing a lever.

Three Perspectives

The dominant view of how animals learn about cause and effect is in terms of associative learning (Dickinson, 1980; Killeen, 1981). Such views typically have neither directed associations, nor intervention in their nomenclature. Although not a principled prohibition, this perspective does not incline consideration of a role for intervention in learning.

Note that views by which the animal is seen as a mere input-output machine, swaying passively with the winds of the world, are obliged to deny a role for intervention in learning. Only if the animal is an independent actor, capable of exerting control over its own actions, might its actions constitute original

causes. Although this prohibition of a role for intervention in learning does not extend to most contemporary views of learning, it is worth bearing in mind.

A secondary, emerging perspective on learning about cause and effect takes inspiration from use of network models of causal reasoning. Notably, Bayesian Networks (Pearl, 2000), initially developed as a set of mathematical tools for dealing with conditional probabilities of systems of interrelated events, provides a basis for a normative theory of causal reasoning, the way a rational agent ought to reason about cause and effect. Having a normative model of how to reason about systems of causal relations provides a top-down aspect to analysis of learning, an endpoint that a learning process ought to ultimately produce. Despite reasonable hesitation concerning normative approaches (Pierce & Ollason, 1987), they have shown their worth (Monteiro, Vasconcelos, & Kacelnik, 2013; Stearns & Schmid-Hempel, 1987). Application of graphical models to human learning has produced promising theories (e.g., Cheng, 1997; Fernbach & Sloman, 2009; Wellen & Danks, 2012). It is a diverse category of theory (Bowers, 2017b). Some such theories focus exclusively on learning, while others say very little about learning. Some focus on statistical relationships, while others stress other sources of information. Some embrace Bayesian Networks, while others distance themselves from it. Such views also vary on the level of analysis at which description in terms of graphical models is intended to apply. Despite these many and striking differences, there is widespread agreement among such views on the value and role of intervention.

Such theories have typically been developed as specifically about human learning, and their successes have so far not carried well to attempts to extend the perspective beyond humans to other animals (Bowers & Timberlake, 2017; Burgess et al., 2012; Dwyer, Starns, & Honey, 2009), despite the benefit of paths already blazed by researchers working with human subjects (e.g., Steyvers, Tenenbaum, Wagenmakers, & Blum, 2003). One aim of the present study is to provide data to stimulate the development of theory that extends the use of graphical models to the learning of other species as well. Our target is not any one theory in particular, but the perspective, and our aim is not especially to rebut, but to incorporate insights from graphical model-based views in a principled way to enrich our understanding of animal learning and cognition.

A third perspective on learning is provided by behavior systems approaches (Domjan, 1994; Timberlake, 1983, 2001; Tinbergen, 1951; overviewed in Bowers, 2017a), by which learning is viewed as a means of fitting systems of interrelated behavior forms to the realities of the specific environment. Bringing intervention into the fold of behavior systems requires very little, just that intervention is an available source of information. It is furthermore fitting because behaving animals actively change their orientation toward system-related stimuli and the motivational context in which these stimuli are encountered in regular ways. If intervention in the course of naturally expressed behavior can guide learning in a manner that helps the animal recover causality from covariation information, providing a potent cue to veridical causality, we expect active learners like rats to use it.

By this view, intervention is expected to exhibit system-specificity, perhaps affecting learning preferentially or differently with specific behaviors and circumstances. Each of the specific actions in a predatory sequence, for instance, can produce relevant outcomes. Given the behavioral specificity under which such outcomes occur, the motivational and environmental contexts will likewise be specific. In this regard, analyses may profit from considering the motivational context in which the animal's behavior produces an outcome, and the significance of that outcome to the evoked system.

All three of these general perspectives can agree that there is value only in associating events between which an apparent covariation accurately indicates a real causal relation (of some kind, perhaps indirect); we had better not associate causally unrelated events (i.e. superstition). All three perspectives can agree that intervention (like precedence and covariation) can provide an important and tangible source of information about causality. All three perspectives can agree that rats have evolved under constraints to behave intelligently under a broad range of complex circumstances. All three perspectives can, therefore, agree that there is compelling motivation for exploring whether and how rats utilize intervention in natural circumstances.

Two Kinds of Causal Association

Intervention permits distinction of two kinds of causal association: those in which A causes B, and those in which covariation of A and B is due to some indirect causal relation between them. Graphical model-based views model the former case as a directed association, depicting an asymmetrical relationship from A to B. In the latter case, one may learn that two events are related, but lack good information about the underlying causality. In the absence of information about intermediate causes, correlations produced by causally indirect associations may be modelled and abbreviated as undirected.

Although the causal relations that may produce undirected correlation come in many forms, which may be both differentiable and important to differentiate for an engineer attempting to model the causal factors relevant to the construction of a bridge, it is unclear whether drawing finer distinctions should be expected of natural cognitive systems. Doing so is hard, and of unclear utility, since most circumstances permit the various underlying causal structures producing a correlation to be treated similarly. Likewise, there are multiple ways for A to cause B (e.g., via an intermediate cause) that can largely be treated similarly. For instance, when you flip a light switch, the action can be treated as causing the light to burn, though the actual causal factors are more complex (hence, our reluctance to abbreviate the focal distinction as direct versus indirect).

In contrast, these two general kinds of association 1) are readily differentiable and 2) warrant different behavior. 1) Cause-effect can be differentiated from other correlations when there are means to intervene on potential causes. Consistently by theories of causal reasoning, intervention gives the actor privileged information about whether the manipulated event is a cause of another or not (Pearl, 2000). 2) A rational actor is expected to perform actions that bring about causes of desired outcomes, but not mere correlates of the same outcomes. Hence, one adjusts the temperature on a thermostat, but not a thermometer. And only if melancholy causes the rain should spring faeries sing sad songs. The relation between a cause and its outcome is qualitatively unlike the relation between other predictors and their correlates; these two general cases warrant different behavior, and are differentiable by intervention.

Despite the apparent primacy of this distinction, and the ambiguous utility of finer distinctions, note that there are, in principle, no undirected associations in Bayesian Networks (Pearl, 2000). This works fine when modelling causality in a deterministic world, assuming that everything has a cause. However, as a model of causal reasoning or causal learning, this presents a problem whenever there are unknown causes. It seems the animal needs to be able to represent correlations for which the underlying causal connection is opaque. Such cases are prevalent and an important part of learning. Undirected edges in a network model foil the otherwise clean Bayesian mathematics. Views that attempt to benefit from that mathematical cleanliness are obliged to model such relationships as mediated by unknown nodes. However, the prohibition in Bayesian Networks is against bidirectional causal links, which is not what is intended here. When we say that the association between A and B is undirected, we do not mean that A and B cause each other, but that they are causally related despite neither causing the other. Hence, speaking in terms of undirected associations presents no principled conflict with Bayesian Networks, although it does deviate from that framework terminologically. This deviation may be counterproductive for engineering aims, yet appropriate for a theory of how animals deal with cause and effect. While it may be rational for an engineer to restrict analyses to directed cause-effect relationships, rats, walruses, and the rest of us appear to need more.

Note that use of intervention as a source of information about the causal relevance of events does not require that learners represent directed and undirected associations qualitatively differently. The distinction can be relevant to learning even by views that draw no such distinction in terms of what learners represent. Different sources of evidence of a relationship among events (e.g., precedence, contiguity, resemblance, intervention) might affect learning qualitatively differently, or only quantitatively. Multiple sources of evidence for a relation might be combined into a single quantitative estimate of association strength and yet learning might still be differently affected by these different sources. For instance, greater experience observing a moderate covariation may produce the same learning as less experience with a stronger kind of evidence, like intervention. Either possibility might

affect learning in quantitative terms, such as speed of acquisition, asymptotic response rate, or resistance to extinction.

The Experiments

If an animal presses a lever to produce a conditional stimulus (CS), three general outcomes are possible, each of which leads the intervening agent to a different conclusion regarding the causal significance of the CS to the unconditional stimulus (US). In the first case, no US appears, or it appears out of synchrony with the CS (unpaired), and no association is formed. In a second case (featured in Experiment 2), the US appears with the CS, and a strong association is formed. Experiments 1 and 3 focus on a third case, in which the CS and US co-vary positively except when the CS is self-produced (e.g., Hyde, 1976). A discernible concurrence between tone and food evinces a real relationship. However, exerting control over the tone does not exert control over food, which is good evidence that the former is not a cause of the latter. The evidence suggests an indirect relationship.

The present experiments are summarized in Table 1. Experiments 1 and 2 feature within-subject designs, using two conditional stimuli. One subject in each pair produces stimulus A by lever contact, presented in both chambers, and her partner’s lever contact similarly produces stimulus B in both chambers. Subjects are tested on both stimuli. The comparison of Experiment 3 is between-subjects; a single stimulus is used, produced by subjects in Group INT, while lever contact by subjects in Group OBS produces no stimulus.

Table 1

<i>Design of Experiments</i>	
Group	Training
Experiment 1 (within-subject; “undirected”)	
INT-A	CSA alone (6/day); Lever produces CSA CSB alone (6/day); CSB (yoked to INT-B partner’s lever contact) CSA + CSB followed by food (12/day)
INT-B	CSA alone (6/day); CSA (yoked to INT-A partner’s lever contact) CSB alone (6/day); Lever produces CSB CSA + CSB followed by food (12/day)
Experiment 2 (within-subject; “directed”)	
INT-A	Lever produces CSA followed by food CSB followed by food (yoked to INT-B partner’s lever contact)
INT-B	CSA followed by food (yoked to INT-A partner’s lever contact) Lever produces CSB followed by food
Experiment 3 (between-subjects; “undirected”)	
Intervene (INT)	CS followed by Food (8/day); Lever produces CS
Observe (OBS)	CS followed by Food (8/day); CS (yoked to INT partner’s lever)

Experiment 1

In Experiment 1, both subjects in each yoked pair produced one stimulus with its own lever contact, and received another stimulus yoked to its partner’s lever contact. This permitted a within-subject comparison, with the aim of heightening sensitivity to hypothesised effects of intervention. Stimuli were presented for only half-second durations to permit tight temporal correspondence between action and

outcome, and reduce interference of one stimulus with the other. Does the ability to produce a stimulus by one's own action change how one learns about its relationship with food?

Methods

Subjects. Eight experimentally naïve female Sprague-Dawley rats served as subjects. For the course of the experiment, they were housed individually on a 12-hr light:dark schedule. They were fed once daily, and maintained at a minimum of 85% of their free-feeding weights. Water was freely available in home cages. Experimental sessions were conducted in the dark portion of the cycle. All animal use was reviewed and approved by the Institutional Animal Care and Use Committee of Indiana University's Bloomington campus.

Apparatus. The apparatus consisted of pairs of chambers made of sheet metal (50 X 35.5 X 30 cm), with a Plexiglas ceiling, and accessed by a hinged Plexiglas front wall. These chambers were in sound-attenuating cabinets in one room, 20 cm apart, and linked to data-logging and control equipment in an adjacent room. The food trough (3.1 cm wide, 5.9cm high, 1.8 cm deep, 4cm above the floor) detected intrusions by an infra-red beam projecting across the entrance. Food delivery consisted of two 45 mg Testdiet purified pellets released into the feeder (one second apart). The retractable lever (3.6cm wide) was on the same wall (5.2 cm from feeder; 4 cm above the floor, 5.1 cm from the back wall). The tone (4 kHz) was emitted from a speaker centred 1.6 cm above the lever. The green light was immediately above the lever; the white light was above the speaker. Three weight-sensitive platforms in the floor of the chambers recorded when depressed (10.2 cm deep from the wall). One was immediately before the food trough (near platform), and two spanned the length of the back wall (far platforms). Gary Lucas' CONMAN software controlled these chambers.

Procedure. Each day, pairs of rat subjects (one from each group) were placed in adjacent chambers. Prior to training, there were two days of exposure to the chamber in which they found just food. On training sessions, subjects received 6 presentations of CSA alone, 6 presentations of CSB alone, and 12 presentations of CSA and CSB together followed by food delivery (2s from onset). For most subjects, CSA was a 0.5-s green jewel light and CSB was a 0.5-s white incandescent light. For one pair of subjects, CSA was a 0.5-s tone, and CSB was a solenoid click from behind the chamber wall. Stimuli were highly discriminable. Presentations appeared at variable intervals, two appearing in every 4-min period (median onset-to-onset spacing: 90 s; minimum: 50s). Lever contact by either subject in each pair would interrupt this progression (except during any CS, or within 10s of a scheduled CS). Lever contact by INT-A subjects produced CSA in its own and yoked chambers; lever contact by INT-B subjects would likewise produce CSB in both chambers. The schedule would then resume from the last 20 s step (enforcing a minimum 30s delay after any lever-produced tone before a scheduled presentation). These sessions (1/day) were of variable duration (median scheduled duration: 45.7min). All subjects completed eight sessions. Two pairs of subjects continued to complete 18 sessions¹.

Measurement and analyses. The conditioning chambers automatically recorded whether, in each 0.1s interval, the subject was: a) in contact with the lever; b) in the food trough; c) on the platform immediately before the food trough; or d) on either of two platforms along the back chamber wall. These counts were summed in 2-s and 10-s intervals, yielding a value of 0-100 for each dependent variable, for each 10s-period measured (and 0-20 for every 2 s). Ten-second periods starting with onset of each scheduled CS were recorded, and the immediately preceding 10-s periods. Repeated-measures analyses of variance (ANOVAs) were conducted on each of the four dependent variables to test effects of four within-subject variables: intervention (INT versus OBS conditions); stimulus type (green light, white light, tone, click); period (10-s period immediately before versus 10-s period from onset of stimulus presentation), and days. Autoregressive covariance structure was assumed in recognition of the potential for progressive changes over the course of training. Given multiple comparisons, decisions about null

hypothesis rejection were based on the Benjamini and Hochberg (1995) procedure for controlling the false discovery rate (FDR = .15).

Results and Discussion

Nosing of the food trough (Figure 1) and presence on the near platform both showed main effects of stimulus period [$F(1, 46.74) = 100.26$, and $F(1, 43.54) = 19.65$, respectively; p 's < .001] and days [$F(17, 132.78) = 4.45$, and $F(17, 128.40) = 4.91$; p 's < .001], indicating learning about the relevance of the CSs to food. No dependent variable showed an effect of intervention [Nosing of feeder: $F(1, 34) = 0.22$, $p = .641$; Near platform: $F(1, 31.49) = 0.42$, $p = .52$; Far platforms: $F(1, 76.12) = 0.29$, $p = .59$; Lever contact: $F(1, 44.16) = 0.01$, $p = .95$].




Figure 1. Nosing of the food trough in 10-second periods immediately prior to CS presentation (pre CS), and beginning from CS onset (onset) for intervened upon stimuli (INT) and control stimuli (OBS) in Experiment 1. Error bars indicate SEM.

Noticing that the first pair of subjects appeared to respond more to the click CS than the tone (estimated difference: 26.00), we used visual CSs for subsequent pairs. Greater conditional nosing of the feeder was observed with the white incandescent light CS than the green jewel light CS (estimated difference: 24.07), suggesting that we again failed to equate the salience of the stimuli [main effect of stimulus quality: $F(3, 32.27) = 7.75$; $p < .001$]. The presence of measurable effects of subtle differences of stimulus salience puts our failure to observe effects of intervention in perspective. By hypothesis, intervention provides a powerful source of information. A small effect of intervention being swamped by subtle effects of salience does not fit expectations about a hypothesised role for intervention in learning.

This intuition was used to establish equivalence bounds (Lakens, 2017) in order to be clearer about whether the data justify acceptance of the null hypothesis, as opposed to mere failure to reject it. Considering an effect of intervention no greater than the estimated effect of stimulus quality as uninterestingly small, we used the least extreme of these two estimates (24.07) to set lower and upper equivalence bounds. The confidence level of the widest confidence interval contained within these bounds (-24.07 to 24.07) is 80.5%. From this, we derived a conservative estimate of the probability of the alternative hypothesis in each tail, given the data [$p(H_1|data) = (1 - .805)/2 = .098$] to compare with the estimated Type 1 error rate [$p(H_0|data) = .641$], to obtain an odds ratio (Rouder, Speckman, Dongchu, Morey, & Iverson, 2009; N.B. These calculations are based on uncorrected, per-comparison error rate estimates; family-wise error rate correction would distort these calculations). The resultant odds ratio [$p(H_0|data) / p(H_1|data)$] is 6.57, indicating that the null hypothesis is 6.57 times more likely given the data than either alternative (greater or lesser). Thus, the data justify the negative conclusion that subjects

nosed the food trough in response to CSs at similar rates whether they could produce the stimulus themselves by lever contact or not.

Lever contact and presence on the far platforms were both periodic and infrequent, producing quite low values in the 10s periods measured; these measures revealed no effect of CS presentation, nor progressive change over days. Although lever contact counts were low in 10-s bins (mean pre-CS and CS periods: 0.33, 0.56, respectively), over the course of the sessions, lever contact was substantial and steady, producing many response-dependent CS presentations (median: 29.25 presentations per session; see Appendix, Figure A1).

In Experiment 1, CSs were short, lasting only 0.5s. This was to maximise the correlation between lever contact and appearance of the response-dependent CS. Since lever contacts during a CS had no additional effect, a shorter CS permitted a higher correspondence between action and outcome. One disadvantage of short CS durations in Experiment 1 was that, for most of the 10-s measurement periods, the CS had already ceased, which may produce a less sensitive test. Thus, in the following experiments, 10-s stimulus durations were used.

One hesitation about Experiment 1 concerns the moderate CS-food correlation it presented. Some views conceive of the quantitative aspect of associations separately from qualitative aspects (e.g., Wellen & Danks, 2012), and by such views, assessment of the directedness of an association is independent of the assessment of its strength. However, intuitively, people are reluctant to attribute direct causation between only moderately correlated events. Perhaps the modest correlation presented in Experiment 1 prevented subjects from attributing direct causation in either condition. To deal with this concern, Experiment 2 maintained a 1:1 relation between each CS and food.

Experiment 2

In the context of clear evidence of learning, the absence of an effect of intervention on conditional responding in Experiment 1 was striking. However, beyond how intervention is expected to influence learning, views differ on what one is expected to learn when presented with covariation information alone, as in the case of the control stimuli (OBS). By some views, learners are expected to readily acquire directed associations between preceding stimuli and reliably following stimuli (e.g., Fernbach & Sloman, 2009). The results of Experiment 1 speak against such views, which would expect stronger response to OBS stimuli than INT stimuli. However, these results are consistent with the contrasting view that OBS subjects would learn an *undirected* association in the absence of good evidence for a directed association, and therefore not differ from INT subjects able to affirm the lack of direct causation by intervening on the CS. By this latter interpretation, if intervention had instead evinced a directed relationship between the stimulus and food, it would outpace learning from concurrence alone. Experiment 2 presented subjects with such a case.

In Experiment 2, one subject could produce CSA by lever contact at any moment in its own chamber and its neighbour's chamber. Meanwhile, its neighbour had similar control over CSB. Unlike in Experiment 1, however, these presentations were followed by food delivery, and so intervening on one CS evinced a directed relationship of that stimulus with food. If rats use intervention to learn about the relationship of each CS with food, learning of a strong CS-food association should be faster for the intervened upon stimulus than the control stimulus.

Methods

Six of the subjects in Experiment 1 (those for whom CSs were lights) served as subjects in Experiment 2 ($N = 6$). The apparatus was identical to Experiment 1, except that food delivery consisted of only one food pellet, and CSA and CSB were auditory stimuli rather than visual. CSs were a 10-s 82-dB white noise and a 10-s 500 Hz tone, emitted from spatially distinct speakers along the same wall, each roughly 14 cm from the lever. Lever contact by INT-A subjects produced CSA followed by food delivery in both yoked chambers, except during a 30-s refractory period beginning with the onset of any CSA

presentation. Lever contact by INT-B subjects similarly produced CSB followed by food in both chambers, except during a similar refractory period. Refractory periods enforced a minimum of 20s between presentations of a given CS, although CSA and CSB could overlap. Sessions continued for 30min, or until a subject produced 40 stimuli. Sessions continued for 12 days, except for one pair of subjects, which were run for only 6 days¹.

Repeated-measures ANOVAs of nosing of the food trough and lever contact were conducted as in Experiment 1. Odds ratios [$p(H_0|data) / p(H_1|data)$] were again used to quantify confidence in equivalence claims, as described above. Given similar measurement, we used the equivalence bounds set in Experiment 1 (-24.07, 24.07 for nosing of the food trough).

Results and Discussion

Main effects of stimulus period [$F(1, 31.34) = 67.73; p < .001$] and days [$F(11, 107.42) = 8.54; p < .001$] for nosing of the food trough indicated learning of the food-predictive quality of the CSs (Figure 2). Again, hypothesis testing revealed no effect of intervention [$F(1, 22.55) = 0.09; p = .76$], nor any interactions with intervention, suggesting similar levels of feeder-directed responding to response-dependent and response-independent CSs. With 93% confidence in the effect being within the equivalence bounds, the odds ratio strongly favoured the null hypothesis (21.77:1).

Note that in Experiment 2, response is measured under two very different conditions. In particular, intervened-upon stimuli are presented specifically when the subject had just exhibited a particular response. Thus, an effect or lack of effect of intervention could be influenced by systematic motivational differences between conditions, a serious limitation of this experiment. However, in the context of there being no effect of intervention, it is meaningful to note that there was a tangible effect of stimulus quality: conditional responding was greater to the white noise than the 500 Hz tone [$F(1, 22.55) = 5.42; p = .029$]. This also did not interact with intervention [$F(1, 21.15) = 0.32; p = .58$], indicating that the effect of stimulus quality was independent of whether the stimulus was self-produced or not (Figure 2). Thus, the absence of any evidence of an effect of intervention is again put in perspective by a small but significant effect of seemingly subtle differences in stimulus salience.




Figure 2. Nosing of the food trough in 10-second periods immediately prior to CS presentation (pre CS), and during both CSs (white noise and 500Hz tone) for intervened upon stimuli (INT) and response-independent stimuli (OBS) in Experiment 2. Error bars indicate SEM.

Lever contact increased over days [$F(11, 93.36) = 2.51; p = .008$; see Appendix, Figure A1]. This was expected given that lever contact was an instrumental response. Lever contact was greater during response-dependent CS presentations than response-independent CS presentations [$F(1, 35.79) = 21.90; p < .001$; cf. Dwyer et al., 2009], which may be regarded as a methodological artefact, since the former were initiated by lever contact, and lever contact occurs in spurts, such that lever contact predicts lever contact. Absence of an interaction of days with stimulus period [$F(11, 92.36) = 0.94; p = .50$] indicates that the increase of lever contact over days was similar during and prior to stimulus presentation.

Experiment 3

Neither of the above experiments encourage the view that rats use intervention as a source of information about causality. However, both experiments featured within-subject designs, which are suited for equating subjects on extraneous variables, including motivational factors. Such factors may be relevant to how rats deal with cause and effect in circumstances typical of the animal's ecological niche. It is worth remembering that the present experiments subject rats to food-related problems, and measure food-related behaviors, and so specific motivational systems are involved (Timberlake, 2002). Experiment 3 uses a between-subjects design to permit comparison of intervening subjects with control subjects that receive the same pattern of CS and US presentation but without the capacity to produce any CS by its own action.

Methods

Ten experimentally naïve female Sprague-Dawley rats served as subjects, similar to those in Experiment 1. The apparatus, measurement, and pre-training exposure were as described above.

For each of 15 days, pairs of rat subjects (one from each group) were placed in adjacent chambers in which they received 8 presentations of a 10-s white noise stimulus (Tone), followed immediately by food delivery. Presentations were scheduled to appear at variable intervals, two in every 7-min period (mean interval: 200s; minimum: 60s). The INT subject in each pair could interrupt this progression at any moment by contacting the lever (except during any Tone, or within 10s of a scheduled Tone). Lever contact by INT subjects produced a 10s Tone in its own and yoked chambers. After these 10s, the schedule would resume from the last 40 s step (enforcing a minimum 50 s delay after any lever-produced Tone before a scheduled presentation). Lever contact by the OBS subject had no such effect. These sessions were of variable duration (median: 32min; interquartile range: 29-36min). An extinction test was conducted on day 16, but those data closely resemble data from the later days of training, and so discussion of that session was omitted from this report as superfluous.

Repeated-measures ANOVAs were conducted on group, and two within-subject variables: period (10-s periods immediately before versus during stimulus presentation), and days (in 7 two-day blocks, i.e., 16 observations in each block). Covariance was modelled as first-order autoregressive, with heterogeneous variance parameters. The Benjamini and Hochberg (1995) procedure was again used as above. Statistical equivalence was assessed for critical null effects using odds ratios [$p(H_0|data) / p(H_1|data)$] by the reasoning and specific equivalence bounds (-24.07, 24.07 for nosing of the food trough) established in Experiment 1.

Results and Discussion

Inferential statistics are given in Table 2. Nosing of the food trough did not differ between groups (Figure 3; $p = .52$; odds ratio = 8.19). Both groups developed conditional nosing of the feeder to the CS (indicated by main effects of period and days, and period by days interaction). The rate of increase was similar between groups (no interaction with group). Post hoc tests confirmed that feeder-directed responding showed no change in pre-CS periods over days. Appendix Figure A2 gives the data of Experiment 3 in greater detail. Lever contact likewise showed no differences between groups, and no change with continued training. Lever contact rates were similar prior to and during the tone.

Table 2

Inferential Statistics from Experiment 3

	Nosing food trough			Lever contact			Near platform			Far platforms		
	F	df	p	F	df	p	F	df	p	F	df	p
Group	0.43	1, 16.46	.524	0.59	1, 16.65	.452	*13.16	1, 14.65	.003	*14.74	1, 17.10	.001
Period	*22.83	1, 27.69	< .001	3.79	1, 21.45	.065	1.32	1, 32.33	.259	0.18	1, 25.24	.676
Days	*5.75	6, 21.28	.001	1.42	6, 33.93	.234	*12.67	6, 40.31	< .001	*3.72	6, 28.38	.007
Period*Group	0.00	1, 27.69	.986	0.22	1, 21.45	.646	0.27	1, 32.33	.604	0.36	1, 25.24	.555
Days * Group	1.29	6, 21.28	.306	0.77	6, 33.93	.597	0.96	6, 40.31	.467	*2.74	6, 28.38	.032
Period * Days	*4.08	6, 22.37	.007	1.23	6, 32.06	.318	0.37	6, 45.96	.897	0.77	6, 29.41	.597
3-way Interact	0.30	6, 22.37	.932	1.07	6, 32.06	.400	0.26	6, 45.96	.954	0.89	6, 29.41	.516

*Null hypothesis rejections, as identified by applying the Benjamini & Hochberg (1995) procedure (FDR = .15)




Figure 3. Nosing of the food trough by both groups, prior to and during presentations, over the 15 days of training of Experiment 3. Error bars indicate SEM.

Location and locomotion. OBS subjects spent more time (pre-stimulus and stimulus periods) on the platform by the feeder than INT subjects (Figure 4a). Both groups increased time by the feeder over days to reach near-ceiling levels (main effect of days; no interaction with group). Pre-stimulus levels appeared to be lower at asymptote for group OBS than group INT, but though this can be brought out by post hoc data prospecting, it was not significant by the planned test. Note that the common strategy of analysing the stimulus — pre-stimulus difference in response levels would have given the impression that INT subjects progressively developed conditional approach, while OBS subjects remained at baseline (drawn in Appendix Figure A3), which appears misleading in light of the untransformed data (Figure 4a).

The pattern of movement away from the food trough, as measured by time on the far platforms, was unsurprisingly largely opposite that of approach (Figure 4b). A group difference was again apparent, particularly in earlier sessions, with OBS subjects rarely and INT subjects frequently on the far platforms. A similar difference was observed both during and prior to stimulus presentation. With training, the high activity of INT subjects dropped to resemble levels seen in group OBS (main effects of group, days, and group x days interaction). The striking drop observed in locomotion over several days of visits to the experimental chambers echoes similar observations described elsewhere (e.g., Bindra & Blond, 1958).

In a general way, the presence of a group difference can be seen as supportive of views that attribute importance to intervention in learning. In particular, OBS subjects were nearer to the food trough than INT subjects, a measure that appears within the interpretive scope of normative views. Furthermore, at least one view (discussed below) expects OBS subjects to learn a stronger association than INT subjects, and so appears to gain support from these results.

However, the same pattern may be interpreted in terms of general search, movement away from the food trough, rather than approach. This interpretation is especially viable given the assumptions of a behavior systems approach, by which it is a motivational complex, involving both focal and general search (as well as consummation), that is conditioned, rather than one response form. Thus, the apparent approach could be primarily driven by either the pull of focal search or the push of general search. Three observations favour the general search interpretation. 1) The group difference is apparent not just following food delivery, but in all measurement periods, including pre-stimulus measurements, as expected of a long-term response, such as general search, but not expected of a shorter-termed response, such as focal search. 2) A more direct measure of focal search, nosing of the food trough, does not differ between groups. 3) The group dissociation of far platform data is tighter than that of the near platform data; the observed pattern of the former is sufficient to explain the observed pattern of the latter, but the converse is not true. In contrast, comparison with Experiment 1, in which subjects behaved more like the OBS subjects of Experiment 3, suggests that it is the presence of a response-independent CS that moves the animal to keep close, rather than the response-dependent stimulus that moves it to wander. Thus, the behavior of OBS subjects appears to be best explained by the pull of focal search, but that of INT subjects, by the push of general search.

The lack of group differences in Experiment 1 also speaks against interpretations by which the far platform activity is due to lever contact irrespective of self-production of a conditional stimulus, since that is present in Experiments 1 and 3. Furthermore, the extent of lever contact did not differ between groups.

One might view the subjects' presence on far platforms as avoidance of the lever. If subjects learn an inhibitory association with the lever (again, given training of the form: Tone+; Lever & Tone-), avoidance of the lever may appear to fit. However, two observations are awkward for this interpretation. First, as training continues, subjects progressively spend more time by the food trough and lever, and less time on the far platforms. This interpretation would expect the opposite. Second, as noted above, subjects continue to press, and both groups do so to similar extents.




Figure 4. **a)** Presence in the vicinity of the feeder (on near platform), and **b)** on the far platforms of both groups, prior to and during presentations, over the 15 days of training. Error bars indicate SEM.

Perhaps INT subjects really are quicker to learn the predictive value of the tone, and that is why they are content to wander the chamber, since they effectively have a 10-s warning with every food delivery. Meanwhile, if OBS subjects are slower to see the tone as a predictor, they may be staying by the feeder because they are not able to anticipate deliveries. Two observations discourage this interpretation. First, INT subjects are on the far platforms more than OBS even once the tone sounds (Figure 4b). Second, with continued training, the behavior of intervening subjects came to resemble that of controls, showing progressively less general locomotion, rather than the converse.

Note that the group difference in general search is visible even in the first day of training. When we look closer, within the first day, we see that groups are similar at the first observations, but quickly diverge over the course of the session (see Appendix Figure A4). The differential development of general search behavior among intervening and observing rats appears to emerge with very little experience. This fits earlier suggestions that learning a connection between lever contact and a resultant stimulus is very rapid, effective even from the first instance (e.g., assumed in Blaisdell et al., 2006). However, note that the circumstance changes after the first scheduled presentation, since this marks the experiment's first food delivery. Thus, the drastic change in behavior appears to occur the moment food is on the scene.

Rats show strong preferences for places, and in a conditioning chamber, a rat often picks a favourite spot to light. Hence, a reasonable thought is that the few subjects selected for condition INT may have just happened to pick spots on the far platforms to wait out the experiment, and the measure does not indicate movement at all, but just individual differences in place preference. However, the group difference in presence on the far platforms was observed before, during, and even after stimulus presentations, at which times INT subjects showed just as much presence in the food trough. At other times, these subjects also worked the lever. These observations make it clear that intervening subjects were not merely ending up farther away between presentations, but were actively moving about the chamber. Meanwhile, OBS subjects were indeed sitting by the food trough. These two groups of animals were behaving qualitatively differently.

The quantity of lever contacts was similar between groups. The number and pattern of CSs and USs presented were identical. What differed between groups was the relationship of lever contact with CS presentation. Intervening upon a predictor of food appears to produce a specific response topography characterised by maintenance of general search activity.

General Discussion

Intervening subjects who pressed a lever for a conditional stimulus and subjects who received the same pattern of stimulus presentations irrespective of their actions showed similar levels of feeder-directed response, a standard criterion of learning. If one were to look no further, this finding would appear to suggest that intervening does not affect learning about self-produced stimuli. However, in a between-subjects analysis, groups did reliably differ in another feeding-related way. While intervening subjects checked the food trough no less, they spent much more of their time away from it; control subjects, in contrast, waited by the food trough for the next delivery. Hence, intervention does change the rat's behavior, but rather than simply enhance or hinder acquisition of a goal-oriented conditional response, the difference appears in behavioral measures that do not have clear interpretation in contemporary theories of causal reasoning.

The group difference of location and locomotion observed in Experiment 3 appeared to be a product of, both, INT subjects moving away from the food trough and OBS subjects keeping near. This divergence of response form suggests that groups learnt differently. By hypothesis, our manipulation involved presenting INT subjects, but not OBS subjects, specific evidence that the CS-US relationship was undirected. As discussed above, not all perspectives dissociate directed from undirected associations, and those that do give no reason to expect rats to respond to undirected associations in feeding contexts with general search especially. We discuss some potentially relevant considerations below in relation to feeding ecology.

Part of this pattern, OBS subjects keeping near to the food source, fits one available hypothesis. It is a frequent assumption (e.g., Blaisdell et al., 2006; Dwyer et al., 2009; Fernbach & Sloman, 2009) that learners will take covariation information alone as evidence of direct causation. Seeing that a stimulus is often followed by food might lead one to interpret that stimulus as a cause of food. By this view, the OBS subject who witnesses only that food differentially follows the stimulus, in the absence of contrary information, should acquire a directed association. In contrast, the INT subject of Experiments 1 and 3 has good information that the same stimulus is not a cause of food, and so acquires an undirected association. By this interpretation, INT subjects should show lower conditional response levels at asymptote than OBS subjects who receive information only about concurrence. The near platform data of Experiment 3, in which OBS subjects were nearer to the food trough than INT subjects, may be regarded as supportive of this hypothesis. However, the group difference appears to be at least partially explained by INT subjects actively moving away from the food trough, which does not fit the mechanistic story of this view. Furthermore, the absence of effects of a more direct measure of goal seeking (nosing of the food trough), and the absence of any such effects in Experiment 1, are awkward for this view.

Does intervention speed learning? If intervention provides an additional, good source of information about causation, whether or not intervening subjects should learn an ultimately stronger CS-food association, or acquire greater asymptotic response levels, an ability to intervene should permit them to learn it faster. There was no clear evidence from any of the above experiments that intervened upon stimuli were learnt about more quickly than control stimuli.

In general terms, the absence of effects of intervention in conditional entry of the feeder is not encouraging of theories of causal reasoning, by which intervention should impart an advantage to learning. However, it is important not to exaggerate the maturity of contemporary theory, and its readiness to be subjected to hard tests to differentiate among clear theoretical alternatives. The notion that intervention is highly relevant to learning is amply motivated. However, this very abstract point, though well grasped in normative terms, is not well represented in theory concerning substantive psychological processes. Hence, the present experiments can provide no empirical basis for exclusion of any current theory of causal cognition. Rather, they should be regarded as primarily exploratory, providing a description of the behavior of intervening and non-intervening rats in a conditioning setting, with the aim of beginning to address questions concerning a potential role for intervention in animal learning.

Predictions of Substantive Views of Causal Cognition

When moving beyond the relative firmness of a normative framework to draw substantive hypotheses about real cognition, one faces the cruel truth that the spareness of normative views leave large gaps that need filling. These gaps in theory are often filled by intuitions, principled or otherwise. The present data permit us to comment on some intuitions that appear in the animal learning literature.

What do OBS subjects learn? Regarding the form of learning, intervention is a point of relative clarity since it permits the learner to dissociate directed associations from other correlations. There is less clarity regarding how an animal should learn about ambiguous cues of causal relatedness, such as concurrence information alone. We argued above against the view that learners take covariation as evidence of direct causation. The opposite view is just as viable. In the absence of clear cues of causation, a subject may initially treat a new association as indicating an undirected relation. If causes are rarer than other correlates, this might be expected. If this is the case, though only the INT subjects will be able to confirm when the association is undirected, OBS subjects will also treat it as such. To assess the strength of an undirected association, the intervening animal has to rely on covariation. Similarly, the control subject has only this kind of information. In this case, the two groups in Experiments 1 and 3 may learn about the stimulus similarly, and so respond similarly. Thus, even assuming a role for intervention, the same pattern might be predicted as assuming no such role. In this case, however, a difference should obtain in Experiment 2, with INT subjects learning a stronger association and showing greater asymptotic responding to the stimulus than OBS subjects. This was not observed.

Do rats discount self-produced presentations? By substantive views of causal reasoning (e.g., Waldmann, Hagmayer, & Blaisdell, 2006), if subjects in Experiments 1 and 3 viewed the intervened-upon stimulus as only indirectly related to food, then lever-produced CS presentations should not have led to expectation of food (*cf.* Blaisdell et al., 2006). Having discounted specifically these non-reinforced presentations, the INT subject receives a strong correlation (effectively 1:1 in Experiment 3) between the CS and food. Meanwhile, the OBS subject receives a number of non-reinforced presentations that it is not able to discount, and so learns a weaker association. If this is the case, INT subjects should more quickly and more strongly acquire conditional responding to the CS. This was not observed, calling into doubt explanations of related experiments that rely heavily on the assumption that rats engage in such discounting (e.g., Blaisdell et al., 2006).

Lever contact. When lever contact produces food in conditioning experiments (as in Experiment 2), rats quickly learn to press, and reliably, pressing rises as training continues (instrumental conditioning). When lever contact produces a stimulus that signals the absence of food, subjects learn to avoid pressing. However, when the lever produces a CS without giving the presser control over food (as in Experiments 1 and 3), rats neither increased nor decreased lever contact with continued training (*cf.* Hyde, 1976). They continued interacting with levers that produced excitatory stimuli, even though lever contact was never rewarded, maintaining initial response rates even after 15 and 18 days of training. As lever contact neither helps nor hinders in this case, it being neither diagnostic of food, nor instrumental, such findings are consistent with all general perspectives considered here.

Views that borrow heavily from associationism might predict, for INT subjects in Experiments 1 and 3, increases in conditional lever contact by secondary reinforcement, or conversely, extinction of lever contact, as a negative discriminative cue or inhibitor, since the lever produces the conditional stimulus in the absence of food (*cf.* Yin, Barnet, & Miller, 1994). Analogously, theories of causal learning may regard the lever as a cause of a cause, or as a causal preventer. This combination of positive and negative influences can fit many patterns, so although the observed maintenance of response is consistent, this is trivial.

Ambiguity remains concerning how substantive views of causal cognition expect learning to differ between the two circumstances of our intervening and non-intervening subjects, in addition to how such learning should manifest in behavior. This permits a wide range of views to claim support from the finding of no group differences in feeder-directed behavior. However, doing so involves making specific commitments. For instance, a theory may claim support from the current results by drawing assertions about the kind of training that justifies attribution of direct causation. This may be a positive step if it instigates more specific claims, which may thereby help resolve existing ambiguities for future work.

These unanswered questions serve as a warning that the conceptual language of causal cognition is still somewhat malleable. Care is warranted to prevent its abuse and misapprehension. For instance, we might have easily told a story about the present data involving a post hoc selection of the above intuitions. Let future theory articulate specific commitments.

An Integrative Approach

A compelling case has been made for the importance of the role of intervention in learning about cause and effect (Pearl, 2000). Previous research on animal learning has focused heavily on statistical relationships among stimuli. Intervention is no less observable, and when available, no less accessible to the animal as a source of information. Able learners like rats may, therefore, be expected to exploit this source of information.

However, typical approaches to causal cognition view behavior in terms of rational acquisition of goals, which limits their interpretive scope to a very narrow subset of behavioral forms. Normative theories are restricted to measures that are readily interpretable in terms of fulfillment of some purpose. What else the animal does while abiding or failing to abide rational norms is beyond the reach of such approaches. In the case of Skinner Boxes and food reward, a normative theory might be able to predict changes in the checking of a food trough, but such views, in principle, provide no basis to interpret

general activity measures. Intervention is predicted to impart an advantage to learning. To the finding that groups differed in movement away from the source of food, such views have nothing to say. Accepting that such behavior matters (Bindra, 1961), theories of this sort are incapable of completeness.

One solution to this general problem lies in integration of complementary approaches. Normative theories bring hypotheses and good reasons for believing them that are not otherwise accessible to many views. However, they do not function well on their own. They do not, alone, provide reasons to expect specific patterns of behavior. They require framing. They have often received anthropomorphic framing, and although we do not wish to appear to disparage anthropomorphism in principle (see Burghardt, 2007), we do wish to stress that normative approaches do not require anthropomorphic framing, and to highlight an alternative.

Within a behavior systems framework, it is meaningful to ask whether and how intervention impacts learning, while taking seriously the conceptual basis of theories that make predictions about this. The present research has isolated patterns of behavior that vary with intervention upon conditional stimuli, but that are not interpretable in terms of current theories of causal reasoning alone. However, embedding analysis of this specific learning circumstance within a behavior systems framework places the results of experiments in an independently interpretable context that provides a more comprehensive view of behavior. Doing so, we find that intervention does appear to affect learning, but not where views that attribute a role to intervention would have looked.

Feeding Ecology and Intervention. The present study reveals an apparent relation between intervention and general locomotion. Subjects who intervened to produce the conditional stimulus by their own action differed from yoked controls only, but strikingly, in location and locomotion. Although the connection between intervention and causality has clear theoretical motivation, a conceptual link with general locomotion has yet to be drawn.

By a behavior systems perspective, general locomotion has clear relevance to feeding among rats. The obtained results, therefore, encourage attention to feeding ecology. In particular, as omnivores, rats show distinct behavior patterns for predatory and non-predatory feeding (Timberlake & Washburne, 1989). If the observed difference in topography of behavior with intervention is related to a corresponding difference in feeding strategy, there may be promise in analysis of the interaction of the conditions under which a rat's feeding behavior produces effects, and the structural regularities in the causal relations involved in a typical predatory circumstance. To gain is a behavioral foothold in the study of causal cognition.

Rats' predatory and non-predatory feeding repertoires may bring different opportunities for intervention. Intervention in the course of predatory and non-predatory behavior sequences may be expected to have different consequences on what the animal learns about self-produced conditional stimuli, and lead to different conditional response topographies. We expect the structure of motivational systems to ready the rat to produce the conditions for active learning in specific contexts.

The intervening subjects showed a strong tendency to move away from the food trough, while non-intervening subjects kept close. Note that these two patterns resemble, respectively, a win-shift pattern, appropriate for distributed foraging scenarios, such as a hunting rat might be expected to encounter, and a win-stay pattern, appropriate for foraging in clumpy food distributions (Olton & Schlosberg, 1978; Timberlake & White, 1990), which might better suit a rat's non-predatory foraging. Note further that these two behavior patterns, which outwardly appear very dissimilar, may be produced by adjustment in the relative timing of search modes. In the parlance of behavior systems, both focal search motivation, moving the animal toward known food sources, and general search motivation, moving the animal elsewhere, have a role in both predatory and non-predatory feeding circumstances. However, they may interact differently in ways that suit different foraging scenarios. Even subtle differences in relative timing of motivational modes can have substantial effects on response topography.

In the present experiments, what differed between conditions was the correspondence of the food-predictive stimulus, not with the reinforcer, but with the animal's own consummatory response. In behavior systems terms, consummation is a third motivational mode, which is tightly related to focal search and general search. Thus, the two distinct response patterns observed varied with the presence or

absence of a temporal relationship between changes in feeding motivation as a whole (i.e., the entire motivational complex) and conditional stimulus presentation.

Consider the following conjecture. By producing the excitatory conditional stimulus in specific consummatory motivational circumstances related to feeding, intervening subjects initially treated the learning circumstance as specifically predatory, and therefore adopted a strategy suited for distributed foraging landscapes, characterized by high levels of general locomotion. In the wild, this would ready the hunting animal to look for the next morsel elsewhere, or track escaping prey. Meanwhile, given stationary stimuli and a single, replenishing food source, non-intervening control subjects rapidly (and intervening subjects eventually) adopted the very different strategy of settling by the food trough, suited for clumpy or replenishing foraging circumstances. The structure of feeding motivation may provide an ontogenetic basis for grounding learning about cause and effect in the ecological niche.

Bringing foraging strategies and the structure of motivational systems into research concerning questions of causal cognition frames analyses in terms of specific ecological contexts, which permits reference to specific selection pressures, a prerequisite for a principled evolutionary approach. Our intention is not to defend any specific hypothesis, but to stress the viability of the approach of bringing analysis of behavior sequences in specific foraging contexts into causal cognition research. Such analyses have only to benefit from alignment with principled normative approaches.

Conclusions

The prediction that an effect of intervention on learning will be visible in goal-directed behavior is neither a minute detail nor an obscure hypothesis, but a central aspect of a large and burgeoning body of theory. Against substantial theoretical pressure for such an effect to appear, it is striking that our attempts to isolate intervention in the somewhat obvious context of lever pressing in rats have failed to find support for these expectations. However, it is not the case that we found no effect of intervention. Affording rats opportunity to produce the food-predictive cue during conditioning appeared to affect general locomotion. The rub is that such measures lack clear interpretation by the very views that motivate the prediction. The ostensible failure these data highlight gives meaning by identifying where the invested views need conceptual help.

With a tip of the hat to Bindra (1961), we note that appeal to normative or computational-level theories limits analyses to a narrow view of behavior, and that these analyses specifically suffer from the exclusion of general activity measures. With a tip of the hat to Tinbergen (1951), we present the beginnings of a positive solution, alignment with behavior systems views, analogous to how such views have been used to provide conceptual bridgework in the past (Timberlake, 1983, 2001). Such views are, in principle, neutral regarding assumptions of contemporary theories of causal reasoning, but permit theoretical access to a broader range of behaviors. By this strategy, we recognize the conceptual basis graphical model-based approaches bring for beginning to deal with intervention, while adopting the very different set of assumptions of a behavior systems approach. A primary strength of behavior systems approaches, provision of a more comprehensive view of behavior, complements a primary weakness of normative approaches, their sparseness. Even if there is agreement on where normative views touch ground, there are broad gaps between these points, perhaps too broad to cross without a guide. The wild claims and malleability in thinking about causal cognition are not about these few, isolated points, but assumptions about how to navigate the space between. A behavior systems model might not yet have these points on its map, but it brings some guidance on the terrain.

Acknowledgements

We gratefully acknowledge the careful reading of and insightful comments on an earlier draft by Jonathon Crystal, Peter Todd, and Colin Allen. Research assistant, Jimmy Mundell contributed helpful data processing assistance. Helpful advice from Indiana University's Statistical Consulting Center convinced us of the importance of strong statistical support. In particular, the help of statistical consultant

Dr. Stephanie Dickinson warrants specific acknowledgement. The manuscript was prepared in part at Psychology Research Centre (PSI/01662), University of Minho, and supported by the Portuguese Foundation for Science and Technology and the Portuguese Ministry of Science, Technology and Higher Education through national funds and co-financed by FEDER through COMPETE2020 under the PT2020 Partnership Agreement (POCI-01-0145-FEDER-007653).

References

- Adams, C. D., & Dickinson, A. (1981). Instrumental responding following reinforcer devaluation. *Quarterly Journal of Experimental Psychology*, *33B*, 109–121.
- Benjamini, Y., & Hochberg, Y. (1995). Controlling the false discovery rate: A practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society, Series B*, *57*, 289–300.
- Bindra, D. (1961). Components of general activity and the analysis of behavior. *Psychological Review*, *68*, 205–215.
- Bindra, D., & Blond, J. (1958). A time-sample method for measuring general activity and its components. *Canadian Journal of Psychology/Revue Canadienne de Psychologie*, *12*, 74–76.
- Bowers, R. I. (2017a). Behavior systems. In J. Vonk & T. K. Shackelford (Eds.), *Encyclopedia of animal cognition and behavior* (pp. 1–8). Cham, Germany: Springer. https://doi.org/10.1007/978-3-319-47829-6_1232-1
- Bowers, R. I. (2017b). Causal reasoning. In T. K. Shackelford & V. A. Weekes-Shackelford (Eds.), *Encyclopedia of evolutionary psychological science*. Cham, Germany: Springer. https://doi.org/10.1007/978-3-319-16999-6_3114-1
- Bowers, R. I., & Timberlake, W. (2017). Do rats learn conditional independence? *Royal Society Open Science*, *4*, 160994. <http://dx.doi.org/10.1098/rsos.160994>
- Blaisdell, A. P., Sawa, K., Leising, K. J., & Waldmann, M. R. (2006). Causal reasoning in rats. *Science*, *311*, 1020–1022.
- Burgess, K. V., Dwyer, D. M., & Honey, R. C. (2012). Re-assessing causal accounts of learnt behavior in rats. *Journal of Experimental Psychology: Animal Behavior Processes*, *38*, 148–156.
- Burghardt, G. M. (2007). Critical anthropomorphism, uncritical anthropocentrism, and naïve nominalism. *Comparative Cognition and Behavior Reviews*, *2*, 136–138.
- Cheng, P. W. (1997). From covariation to causation: A causal power theory. *Psychological Review*, *104*, 367–405.
- Dickinson, A. (1980). *Contemporary animal learning theory*. Cambridge: Cambridge University Press.
- Dickinson, A. (2001). Causal learning: Association versus computation. *Current Directions in Psychological Science*, *10*, 127–132.
- Domjan, M. P. (1994). Formulation of a behavior system for sexual conditioning. *Psychonomic Bulletin & Review*, *1*, 421–428.
- Dwyer, D. M., Starns, J., & Honey, R. C. (2009). "Causal reasoning" in rats: A re-appraisal. *Journal of Experimental Psychology: Animal Behavior Processes*, *35*, 578–586.
- Fernbach, P. M., & Sloman, S. A. (2009). Causal learning with local computations. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *35*, 678–693.
- Hyde, T. S. (1976). The effect of Pavlovian stimuli on the acquisition of a new response. *Learning and Motivation*, *7*, 223–239.
- Killeen, P.R. (1981). Learning as causal inference. *Quantitative Analyses of Behavior*, *1*, 89–112.
- Lakens, D. (2017). Equivalence tests: A practical primer for t tests, correlations, and meta-analyses. *Social Psychological and Personality Science*, *8*, 355–362.
- Leising, K. J., Wong, J., Waldmann, M. R., & Blaisdell, A. P. (2008). The special status of actions in causal reasoning in rats. *Journal of Experimental Psychology: General*, *137*, 514–527.
- Monteiro, T., Vasconcelos, M., & Kacelnik, A. (2013). Starlings uphold principles of economic rationality for delay and probability of reward. *Proceedings of the Royal Society-B*, *280*, 1–6.
- Olton, D. S., & Schlosberg, P. (1978). Food-searching strategies in young rats: Win-shift predominates over win-stay. *Journal of Comparative and Physiological Psychology*, *92*, 609–618.
- Pearl, J. (2000). *Causality*. London: Cambridge University Press.
- Pierce, G. J., & Ollason, J. G. (1987). Eight reasons why optimal foraging theory is a complete waste of time. *Oikos*, *49*, 111–117.
- Polack, C. W., McConnell, B. L., & Miller, R. R. (2013). Associative foundation of causal learning in rats. *Learning & Behavior*, *41*, 25–41.

- Rouder, J., Speckman, P., Dongchu, S., Morey, R., & Iverson, G. (2009). Bayesian t-tests for accepting and rejecting the null hypothesis. *Psychonomic Bulletin & Review*, *16*, 225–237.
- Shanks, D. R., & Dickinson, A. (1991). Instrumental judgment and performance under variations in action-outcome contingency and contiguity. *Memory & Cognition*, *19*, 353–360.
- Stearns, S. C., & Schmid-Hempel, P. (1987). Evolutionary insights should not be wasted. *Oikos*, *49*, 118–125.
- Steyvers, M., Tenenbaum, J. B., Wagenmakers, E.-J., & Blum, B. (2003). Inferring causal networks from observations and interventions. *Cognitive Science*, *27*, 453–489.
- Thorndike, E. L. (1911). *Animal intelligence*. New York: Macmillan.
- Timberlake, W. (1983). The functional organization of appetitive behavior: Behavior systems and learning. In M. Zeiler & P. Harzem (Eds.), *Advances in analysis of behaviour*, Vol. 3 (pp. 177-221). London: John Wiley & Sons.
- Timberlake, W. (2001). Motivational modes in behavior systems. In R. R. Mowrer & S. B. Klein (Eds.), *Handbook of contemporary learning theories* (pp. 155–209). Mahwah, NJ: Erlbaum.
- Timberlake, W. (2002). Niche-related learning in laboratory paradigms: The case of maze behavior in Norway rats. *Behavioural Brain Research*, *134*, 355–374.
- Timberlake, W., & Washburne, D. L. (1989). Feeding ecology and laboratory predatory behavior toward live and artificial moving prey in seven rodent species. *Animal Learning & Behavior*, *17*, 1–10.
- Timberlake, W., & White, W. (1990). Winning isn't everything: Rats need only food deprivation and not food reward to efficiently traverse a radial arm maze. *Learning and Motivation*, *21*, 153–163.
- Tinbergen, N. (1951). *The study of instinct*. London: Oxford University Press.
- Waldmann, M. R., Hagmayer, Y., & Blaisdell, A. P. (2006). Beyond the information given: Causal models in learning and reasoning. *Current Directions in Psychological Science*, *15*, 307–311.
- Wellen, S., & Danks, D. (2012). Learning causal structure through local prediction-error learning. In N. Miyake, D. Peebles, & R. P. Cooper (Eds.), *Proceedings of the annual meeting of the cognitive science society* (pp. 2529–2534). Austin, TX: Cognitive Science Society.
- Yin, H., Barnet, R. C., & Miller, R. R. (1994). Second-order conditioning and Pavlovian conditioned inhibition: Operational similarities and differences. *Journal of Experimental Psychology: Animal Behavior Processes*, *20*, 419–428.

1. The reported experiments compose an abandoned series of exploratory pilot experiments. The small procedural shifts throughout are simply our explorations of parameter space. Our aim had been to isolate an effect of intervention from which to begin an analysis. We worried that our failure to find such an effect might have been due to specific design choices (e.g., the relative salience of stimuli), so we changed them (and got the same results). Only later did we realise that our consistent failure to find such an effect was itself meaningful. We apologize for some messiness.

APPENDIX

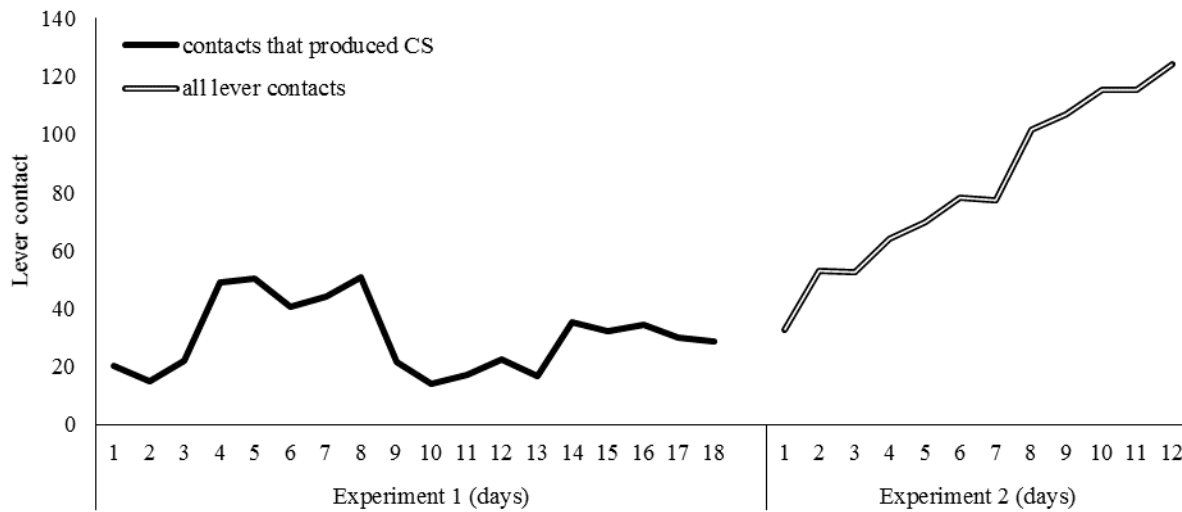


Figure A1. Lever contact over sessions from Experiments 1 and 2. Although presented together here for comparison, the data of these two experiments are not *directly* comparable for several reasons. First, the sessions were notably longer in duration in Experiment 1 than 2. Second, for Experiment 1, these data are actually just showing the number of self-produced presentations, and so only a conservative estimate of the lever contact of only Group INT subjects.




Figure A2. Responding in Experiment 3, prior to, during, and after moments of conditional stimulus presentation (10s periods) in 2s periods over the 15 days of training and ultimate test session. Dotted lines indicate Group INT; solid lines indicate Group OBS.




Figure A3. The stimulus — pre-stimulus difference in presence on the near platform in Experiment 3 over days.

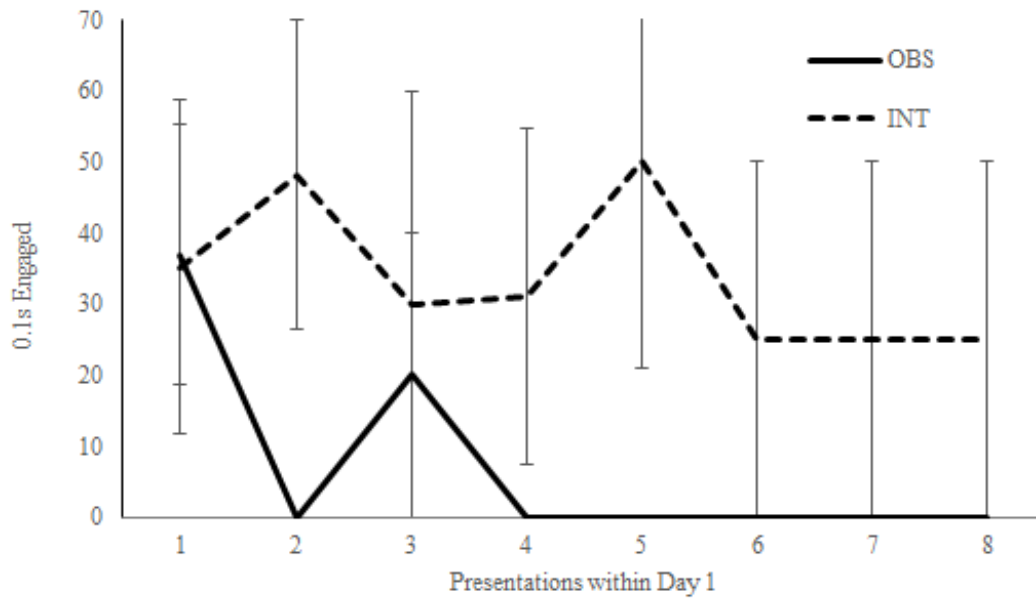


Figure A4. Within the first day of Experiment 3, presence on the far platforms immediately prior to the eight response-independent stimulus presentations, in order. There were insufficient data within the first day to ascertain whether the apparent interaction was reliable or not. Error bars indicate SEM.