

Contingency and Contiguity Trade-Offs in Causal Induction

Marc J. Buehner and Stuart J. McGregor
Cardiff University, United Kingdom

Five experiments investigated the roles of contingency and temporal contiguity in causal reasoning, and the trade-off between them. Participants observed an ongoing, continuous stream of events, which was not segmented into discrete learning trials. Four potential candidate causes competed for explanatory strength with respect to a single dichotomous effect. The effect was contingent on two of these causes, with one of these (A) having a higher probability of producing the effect compared to the other (B), while B was more contiguous to the effect than A. When asked to identify the strongest cause of the effect, participants consistently and reliably selected A, as long as it was not separated from the effect by more than 2.5s. The extent of preference diminished, however, as the contiguity gradient between A and B increased. Beyond 2.5s, the high-probability, but low-contiguity cause A was seen as equally strong as the low-probability, but high-contiguity cause B, and both reliably stood out compared to the remaining two non-contingent distracter items. This apparent trade-off between contingency and contiguity, rooted in contrasting two of David Hume's (1739/1888) fundamental cues to causality, has important implications for psychological and statistical models of causal discovery, learning theory, and artificial intelligence.

How do humans and other intelligent systems learn that one thing causes another? The contemporary cognitive science approach to this problem of induction can be traced back to David Hume (1739/1888), who famously argued that our sensory system is not equipped to directly perceive causality. Instead, he argued, reasoners have to interpret sensory experiences to create a mental representation of causality. Hume identified three principles underlying the formation of causal impressions: i) temporal priority of the cause c before the effect e , ii) temporal and spatial contiguity between c and e , and iii) constant conjunction between c and e . Only the latter two principles are of relevance to cognitive scientists, as the need for temporal priority of c over e is usually not debated (Reichenbach, 1956; but see Savastano & Miller, 1998 for discussion of bi-directional associations).

Computational approaches of causal induction have almost exclusively focused on the third Humean principle, which is commonly referred to as cause-effect *contingency*. Just how exactly contingency gives rise to causal impressions is still subject of a hot debate in the field. Suggestions range from using contingency (ΔP) - calculated by the difference between the two conditional probabilities: $P(e|c) - P(e|\neg c)$ - as a direct measure of causal strength (Allan & Jenkins, 1980; Jenkins & Ward, 1965) to various judgment rules (e.g., Anderson & Sheu, 1995; Mandel & Lehman, 1998; White, 2003). An alternative suggestion (Shanks & Dickinson, 1987) is that causal learning may be no different from associative learning as exemplified by Rescorla & Wagner's (1972) model of Pavlovian conditioning. More recently, however, Cheng (1997) showed that all of

Part of this work has been presented at the 27th Annual Conference of the Cognitive Science Society in Stresa, Italy. This article was written while MJB was on study leave at the Max Planck Institute for Evolutionary Anthropology, Leipzig, Germany, and at the Max Planck Institute for Psychological Research, Munich, Germany. We thank Jacky Boivin and Todd Bailey for helpful discussion of statistical methods to analyze choice data, and Mike Oaksford and Ulrike Hahn for comments on an earlier draft. Correspondence concerning this article should be addressed to Marc Buehner, School of Psychology, Cardiff University, Cardiff, CF10 3 AT, Wales, UK. (BuehnerM@Cardiff.ac.uk).

the above approaches fall short of representing causality as an unbound variable (Holyoak & Hummel, 2000), and suggested a computational causal power approach. A related approach (cf. Buehner & Cheng, 2005) has been to model causal induction as Bayesian inference (e.g. Steyvers, Tenenbaum, Wagenmakers, & Blum, 2003; Tenenbaum & Griffiths, 2001).

The vigorous debate over the computational details of covariation assessment resulted in a lack of integration with the second Humean cue – contiguity (but see Young, 1997). The majority of recent experimental studies have investigated how variations in contingency influence causal assessment; contiguity was never manipulated in these studies, and was usually kept at an immediate level. Earlier work on causal reasoning, however, often focused on contiguity. Michotte (1946/1963) observed that even very short delays render an illusion of causal launching non-causal. In a completely different domain, Shanks, Pearson and Dickinson (1989) reported that people fail to distinguish causal from non-causal actions in an instrumental learning task when the action-outcome delay exceeded two seconds. The importance of Hume’s second principle was acknowledged in early, non-computational psychological theories of causal induction (Einhorn & Hogarth, 1986; Young, 1995).

A Computational Analysis of the Role of Contiguity in Causal Induction

Hume’s analysis of causal induction is very compelling, and the importance of both contiguity and contingency has strong intuitive appeal. What exactly is the role of contiguity in causal induction, though? If one wanted to build an intelligent system to learn about causal relations in its environment, what role would one want to bestow upon contiguity for the system to be optimally adaptive? There are two potential answers to this problem, and the goal of this paper is to shed light on the relation between them.

The Moderating or Mediating Role of Contiguity

One straightforward answer would be that contiguity in itself should play at best a very subordinate role. Intelligent systems learn about causal relations in order to *predict* and *control* their environment (cf. Cheng, 1997). Such abilities arguably exploit information about *regularity* or *contingency*: knowing to what extent one can infer that *c* is followed by *e*, or that acting out or producing *c* is a good way to achieve *e* (cf. also Causal Bayes Nets models such as Griffiths & Tenenbaum, 2005; Pearl, 2000). Contiguity’s role from such a perspective is mostly limited to facilitating or enabling causal discovery. It is evident that contiguous regularities are easier to detect than delayed ones: time-lagged regularities are harder to notice, because event information needs to be kept in memory for longer; as the cause-effect interval increases, the number of potentially intervening (alternative) causes that need to be taken into account increases. In

short, identification of causal relations becomes increasingly difficult as contiguity decreases. And indeed, experimental evidence (e.g. Buehner & May, 2003; Shanks et al., 1989) has demonstrated that – everything else being equal – people have a harder time detecting delayed compared to immediate causal relations.

Contiguity as a Dominating Principle

Sometimes information contained in cause-effect timing can be so salient that it overwrites other important information. More specifically, cause-effect contiguity has been reported to be sufficient to create impressions of causality, despite an absence of contingency (Siegler & Liebert, 1974; Mendelson & Shultz, 1976; Shultz, 1982). Schlotmann (1999) studied how well young children learn the parameters associated with delayed and immediate causal mechanisms inside a toy. She reported that 5-7 year olds consistently preferred a contiguous, immediate cause, over a delayed alternative, even when this choice openly conflicted with well-established knowledge that a delayed mechanism was at work in the toy.

More recently, however, Buehner and McGregor (2006) obtained results that challenged a universal contiguity bias. Participants were presented with an apparatus that was a priori more representative of a delayed than an immediate causal relation. When presented with both short and long delay causal relations, they consistently rated long delay conditions as strongly causal, while short relations were sometimes judged to be non-causal. Buehner and McGregor (2006) argued that *canonical* timeframes, prototypical of a given scenario or mechanism, may overshadow learning of new, non-canonical timeframes. The nature and extent of timeframe biases thus would be intricately linked to the specific scenario, and what the reasoner already knows about it.

In everything described so far, contingency has taken a lead role: Knowing *whether* is more important than knowing *when*. Knowing *when*, according to the above analysis merely helps to know *whether*, but serves no other purpose beyond that. In other words, temporal information would not influence assessments of causal strength in any principled way beyond its potential interference with the ability to detect and assess covariations accurately. More precisely, an intelligent system with no limits on memory or attention should – in the absence of prior knowledge about plausible timeframes – judge delayed regularities just as causal as immediate ones under this analysis.

Contiguity as an Independent Carrier of Causal Information

It is also possible to assign Contiguity a more fundamental role beyond modulating, mediating, or interfering with causal inference based on contingency information. Contiguity itself could carry causal significance. If a patient is terminally ill with a life expectancy of only three months, an ill-judged intervention followed by his immediate death would still be seen as directly responsible for it, even though death would have occurred anyway, albeit at a later time. Likewise, we routinely use medicines that bring about instant relief of a

symptom that would dissipate on its own anyway (e.g. headache), because we consider the speeded-up change causally effective and desirable.¹

Greville and Buehner (2007) recently reported evidence in support for an interaction between contiguity and contingency in causal judgment. They presented participants with tabular data providing information about cause-effect contingencies. Each individual case was marked as having either shown the effect under investigation, or as not having shown the effect. In addition to this frequency information, the tables also contained information about *when* within a critical five-day period each effect had occurred. Thus, the temporal distribution of effects could vary, while contingency (when computed across the total timeframe) could be held constant. Greville and Buehner (2007) showed that participants took both contingency and contiguity into account when making causal judgments. More specifically, identical contingencies were judged as more causal when the temporal distribution peaked near the administration of the cause than when it peaked far from it. In addition, zero-contingencies were interpreted as generative or preventive, when the contiguity was high or low, respectively. Qualitative reports from participants indicated that participants took the entire contingency into account, rather than calculating contingency only over a truncated subset of time, say the first two days. These subjective reports also revealed that participants considered temporal distribution in their judgments. Contingency and contiguity thus acted in concert to influence causal judgments. A similar sensitivity to both contingency and contiguity in real-time instrumental learning paradigms has been reported by Anderson and Sheu (1995) and Wasserman and Neunaber (1986).

Contingency vs Contiguity

Our analysis has outlined two potential roles of timing in causal induction. Firstly, timing (in concert with assumptions about it) can determine how continuous evidence is parsed into discrete units feeding into contingency-based models of causal induction. As such, contiguity serves to *select* potential causal candidates, which then are evaluated based on their statistical relation to the effect in question. Secondly, time can exert a direct influence on causal inference in continuous time. The question then is whether, and how, time and contingency are traded off against one another. This is especially pertinent if one considers multiple competing causes. With multiple options, it can easily be that one has a high value on one dimension, but a low value on the other, while another competitor has exactly the opposite value arrangement. Consider a choice between two varieties of crops, where one produces with a very high probability a yield per year, and

¹Naturally, the causal pathways by which the outcome (e.g. termination of headache) is effected differs between the medication present and absent cases, but – as discussed above – we usually cannot directly observe the causal pathway (e.g. whether the salicylic acid prevents the release of prostaglandins). In the absence of an observable difference in mechanism, we are left simply with information about the presence and absence of candidate causes, and information about whether and when the effect occurred.

another produces with moderate probability a yield per quarter. Which of the two crops is deemed more successful at producing a yield clearly depends on whether we put more emphasis on probability of yield, or on time of delivery. Is there an a priori preference for one or the other? Early developmental work (Mendelson & Shultz, 1976; Shultz, 1982; Siegler & Liebert, 1974) tried to pit contingency and contiguity against one another, but the results were somewhat inconclusive, as manipulations were often confounded with knowledge of mechanism. Mendelson and Shultz (1976), for instance, reported that whether a non-contiguous but contingent cause was preferred over a contiguous but non-contingent cause depended on variations in the physical make-up of the apparatus, and whether these variations were commensurate with the experienced delay. Ideally, one would want to study contingency-contiguity trade-offs in situations where a causal mechanism with a neutral timeframe allows both contiguous and delayed causal relations. Only if both cues can be evaluated neutrally and without any prior hypotheses about specific values, is it possible to assess the true trade-off between them.

The remainder of this paper will investigate which cue carries more weight in situations where two causes with diametrical values of contingency and contiguity compete against each other in a scenario free of a priori timeframe assumptions.

Experiment 1

We developed a new experimental methodology aimed at studying the trade-off between contingency and contiguity in causal induction. We adopted Mendelson and Shultz's (1976) idea to pitch two causes, each with a high value on one, but a low value on the other dimension against each other. More specifically, one cause (A) had a higher contingency with respect to the effect than the other (B), but at the same time A was less contiguous with the effect than B. Unlike in Mendelson and Shultz' study, however, A and B were fully independent of each other and there was no interactive causal influence (Novick & Cheng, 2004) beyond the individual causal strengths. Furthermore, it was important to couch the task in a novel context so that participants would not have any pre-conceived notions of mechanism or expectations of time-frames. This allowed us to study contingency-contiguity trade-offs in a purely bottom-up manner.

To this end, we created a "Stargate" scenario: Participants were told they would observe a group of UFOs orbiting around a stargate; each UFO would attempt to open the gate. Because each UFO would use a unique signaling technique, some would be more successful than others at opening the gate, and some could open the gate faster (if successful) than others. Participants' task was to determine which UFO was most successful at opening the gate.

In designing the task, it was essential to avoid a discretely marked trial structure. In an observational causal learning task involving delays, discrete trial boundaries remove the event-parsing aspect of causal learning (see Allan, Tangen, Wood, & Shah, 2003). Allan et al.'s study adapted the design of Buehner and

May's (2002) but participants were presented with individual learning trials on which the presence or absence of the cause was paired with the presence or absence of an immediate or delayed effect. In other words, participants no longer had to decide whether cause and effect co-occurred on a given occasion; this information was provided through the trial boundaries. All they had to do was to decide whether the observed pairing with a given degree of temporal contiguity was or was not representative of the hypothesized causal relation. It is evident that such artificial judgment tasks bear little resemblance with causal discovery.

Our task presented participants with continuous event streams that were not clearly divided into individual learning trials. Although the event stream was controlled by an underlying trial structure, the appearance to the participant was one of a continuous sequence of events. We strongly encourage readers to watch sample stimuli provided at <http://www.cf.ac.uk/psych/home2/buehner/stimuli/>.

Method

Participants

Ninety-nine undergraduate students from Cardiff University participated to fulfill part of a course requirement.

Apparatus and Procedure

Event sequences were programmed using Macromedia Director, and displayed on a computer screen. The displays represented a 'stargate' in the middle of the screen, which was either open or closed, and four static UFOs, arranged near each corner of the gate. Each UFO had a unique color scheme for its two windows. A 'signalling' UFO was displayed with an overlay of colored stripes, with the color pattern matching the color scheme of the windows. By default, the stargate was closed, and UFOs were inactive. Activity (open gate, active UFO) was scheduled by the program and lasted 500ms. Figure 1 displays sample stimulus materials.



Figure 1. Stimulus Materials. The top right UFO is “active” in this screenshot, and the gate is closed.











Each participant worked on four conditions presented in random order, each consisting of a sequence of trials, presented in random order. The appearance to the participant was one of a continuous stream of events; trial delineation was not made explicit to the participant, and trials were not marked. In each condition, activity of two UFOs (A&B) was probabilistically linked to the opening of the gate, with A always having a higher probability of opening the gate than B as specified in the Design section. Activity in the other two UFOs (C&D) was unrelated to the gate opening. The gate only opened conditional on activity in A or B; in other words the base-rate of the gate opening was zero. The locations and color schemes of each of the four UFOs were randomized for each condition.

The event-structure was organized as follows: If A or B was scheduled to be active on a given trial, they would emit a signal at a random point during the first 5 seconds of that trial. If the signaling was successful, the gate opened for 500ms after the relevant delay. The length of trials was determined by the delay associated with A, and was set to 6.67 seconds + the delay associated with A. Activity in A and B was independent of each other, so that on some trials both A and B would signal; on such trials, A and B would both produce the effect according to their respective probability and delay. In other words, A and B were truly independent of each other and did not interact to produce the effect (Novick & Cheng, 2004). Table 1 lists the various types of event combinations that were necessary to implement five probability gradients (see Design section). For each trial, there was a 25% chance that C and/or D would activate, at a random time throughout the trial. This happened entirely independently of activity in A or B, or the occurrence of the effect. In other words, C and D could occur at any time throughout the experiment, independently of A and B, or the effect.

Participants observed the event streams for each condition, and then had to indicate which UFO was most successful in opening the gate. To this end, they were told to imagine they could “zap” one of the UFOs to emit a signal, and were asked to decide which UFO they would zap in order to open the gate. The experiment lasted about 40 minutes.










Table 1

Number of trial types occurring on each of the five different between-subject probability gradients.

Cond.	P_A	P_B										
a	.75	.25	12	4	4	12	3	9	1	3	16	
b	.75	.50	12	4	8	8	4	8	2	2	16	
c	1.0	.25	16	0	4	12	4	12	0	0	16	
d	1.0	.50	16	0	8	8	8	8	0	0	16	
e	1.0	.75	16	0	12	4	12	4	0	0	16	

Note:

P_A and P_B : Probability that Cause A or B are followed by the Effect.

-  Cause a and Effect associated with a
-  Cause a and no Effect
-  Causes a and b and Effects associated with a and b
-  Causes a and b and Effect associated with a only
-  Causes a and b and Effect associated with b only
-  Cause b and Effect associated with b
-  Cause b and no Effect
-  Causes a and b and no Effect
-  No Causes, No Effects.

Design

The two variables of interest, Contingency and Contiguity, were controlled as follows. A always opened the gate with a higher probability than B. The extent of the probability gradient between A and B was manipulated between participants in five conditions: a) 75% vs. 25%; b) 75% vs. 50%; c) 100% vs. 25%; d) 100% vs. 50%; and e) 100% vs. 75%. Participants were randomly allocated to one of these pairs of probabilities.

The cause-effect contiguity of the low-probability cause (B) was always set to 500ms. The contiguity of the high-probability cause (A) varied within participants across four conditions, and took values of 500ms, 1000ms, 1500ms, and 2000ms. Consequently, trial lengths for these conditions were 7.17s, 7.67s, 8.17s, and 8.67s. 64 learning trials were joined seamlessly per condition. Table 1 lists how many trials of each type occurred in each of the five probability gradients.

Results

Table 2 lists the percentages of choices for the high (A) and low (B) probability cause, as well as the two non-causal distracter items (C) across the five probability gradients, and the four levels of contiguity of A. Inspection of Table 2 suggests that participants consistently selected the high-probability cause A, even when its contiguity with respect to the effect was degraded. This preference seems to be stable across the various probability gradients we implemented.

Table 2

Experiment 1: Percentages of choice for the high-probability cause (A), the low probability cause (B), or one of the two non-causal events (C), broken down by between-subject probability gradient (across) and within-subject contiguity-gradient for A (down). P_A and P_B refer to the probabilities that A or B is followed by the effect.

		P_A	0.75	0.75	1	1	1
		P_B	0.25	0.5	0.25	0.5	0.75
A	500ms		0.80	0.85	1.00	0.84	0.80
	1000ms		0.80	0.70	0.90	0.84	0.70
	1500ms		0.60	0.70	0.95	0.68	0.55
	2000ms		0.60	0.65	0.80	0.63	0.70
B	500ms		0.05	0.15	0.00	0.11	0.15
	1000ms		0.10	0.30	0.10	0.16	0.25
	1500ms		0.30	0.25	0.00	0.21	0.40
	2000ms		0.30	0.30	0.20	0.32	0.30
C	500ms		0.15	0.00	0.00	0.05	0.05
	1000ms		0.10	0.00	0.00	0.00	0.05
	1500ms		0.10	0.05	0.05	0.11	0.05
	2000ms		0.10	0.05	0.00	0.05	0.00

There is no standard statistical procedure to analyze polytomous choice data in a mixed within-between subjects design as we have implemented in this experiment. We thus will not test for interactions between the contiguity and contingency gradients. Because it is not theoretically interesting to investigate the effect of contiguity on each level of probability gradient separately, we decided to

pool the data across the five probability gradients and thus obtain a larger sample and more statistical power. All statistical analyses are based on an alpha-level of 0.05 with Bonferroni-corrections for multiple tests, where applicable.

Figure 2 plots the percentages of choices for A, B, and C, collapsed over the probability gradients. Again, it can be seen that the preference for the high-contingency cause (A) is robust; although it appears that the degree of preference of A over B diminishes as the contiguity of A decreases. Choices for the unrelated distracter causes (C&D) were below 20% in all conditions, and thus well below the chance level of 50%, (all $ps < 0.001$ on a Binomial test).

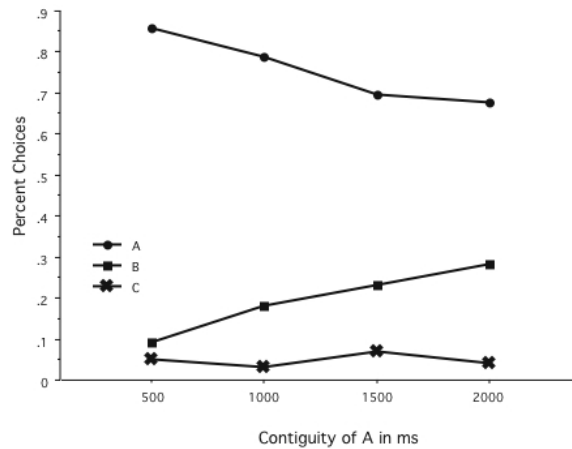


Figure 2. Experiment 1: Percentage of Participants (N=99) choosing Causes A or B or one of the two unrelated events (C&D).

In order to analyze the influence of contiguity on the preference for the high-contingency cause, we constructed three separate and mutually exclusive dichotomous choice variables for causes A and B, and the two unrelated events (C&D) as follows. If a participant selected cause A in a given condition, the A-choice variable was given a value of 1, while the B-choice and C-choice variables were assigned a value of 0. Thus, each participant would score one entry of 1 and two entries of 0 in each condition. The proportion of choices for A was significantly higher than that for B across all four levels of contiguity (all $ps < 0.001$ by sign test). Cochran's Q tests with corrected alpha-level ($p = 0.017$) were conducted to assess the influence of contiguity on choice patterns. Choices for A were significantly affected by variations in A's contiguity with respect to the effect, $Q(3,99) = 13.54$; inspection of Figure 2 suggests that this influence was such that when A's contiguity with the effect decreased, the frequency of A-choices decreased accordingly. Choices for B, on the other hand appear to increase as A's contiguity decreased, $Q(3,99) = 14.42$. Choices for the two unrelated causes were not affected by variations in A's contiguity, $Q(3,99) = 2.56$.

Discussion

Participants were clearly able to extract contingency information from a continuous stream of events that contained no observable trial boundaries. They reliably and consistently identified the cause that was followed by the effect with the highest probability among a choice of four. Moreover, this preference for a high-probability cause was maintained in the face of degraded contiguity: although B was highly contiguous with the effect (500ms), A was consistently preferred as the stronger and more effective cause due to its higher probability, even when A was separated from the effect by as much as 2s, and even when the probabilistic contrast between A and B was as small as 0.25. The extent of this overall preference decreased, however, as the contiguity contrast between A and B increased. Experiment 1 thus suggests that people put more importance on contingency than on contiguity as reliable cues towards causality; at the same time, there seems to be some trade-off between the two, with participants shifting more weight on contiguity, as the contiguity contrast increases.

Experiment 2

Shanks et al. (1989) reported that participants failed to distinguish causal from non-causal actions when the action-outcome interval exceeded two seconds (but see Buehner & May, 2003). Perhaps the contiguity gradient in Experiment 1 was not steep enough to observe a shift from contingency to contiguity. Experiment 2 thus replicated Experiment 1, but employed a larger contiguity contrast between causes A and B: while B was still associated with a 500ms delay, A's delay could take on values of 2500ms, 3250ms, and 4000ms.

Method

Participants

Sixty-five undergraduate students from Cardiff University participated to fulfill part of a course requirement.

Apparatus, Procedure, and Design

The same design and procedure as in Experiment 1 was employed, except that the contiguity of A could take on values of 2500ms, 3250ms, and 4000ms (varied within-subjects), while the contiguity of B remained at 500ms. Trial length was still determined as a function of the delay of A. Consequently, trials were 9.17s, 9.92s, and 10.67s long. The same five probability gradients as in Experiment 1 were varied between subjects.

Results and Discussion

Table 3 lists the percentages of choices for the high (A) and low (B) probability cause, as well as the two non-causal distracter items (C) across the five probability gradients, and the three levels of contiguity for A. Inspection of Table 3

reveals that A is preferred over B across all probability gradients, when A's contiguity with the effect is 2500ms. This pattern breaks down when A's contiguity is degraded to 3250ms, although A is still somewhat preferred in conditions with a comparatively large probability contrast (0.50). With a 4s delay, A loses its attractiveness even more, although it can still maintain it in the 1.0 - 0.25 condition.

Table 3

Experiment 2: Percentages of choice for the high-probability cause (A), the low probability cause (B), or one of the two non-causal events (C), broken down by between-subject probability gradient (across) and within-subject contiguity-gradient for A (down). P_A and P_B refer to the probabilities that A or B is followed by the effect.

		P_A	0.75	0.75	1	1	1
		P_B	0.25	0.5	0.25	0.5	0.75
A	2500ms		0.77	0.62	0.92	0.69	0.69
	3250ms		0.54	0.38	0.77	0.46	0.46
	4000ms		0.23	0.31	0.77	0.54	0.69
B	2500ms		0.15	0.38	0.00	0.31	0.31
	3250ms		0.31	0.54	0.15	0.54	0.54
	4000ms		0.62	0.62	0.15	0.46	0.23
C	2500ms		0.00	0.00	0.08	0.00	0.00
	3250ms		0.15	0.08	0.00	0.00	0.00
	4000ms		0.08	0.08	0.08	0.00	0.08

Figure 3 plots the percentages of choices for A, B, and C, collapsed over the probability gradients. The overall impression is that A is no longer preferred over B when A's delay exceeds 2.5s. The preference pattern does not reverse, however. Overall, B is never picked more often than A, but A and B appear to be equally attractive when A's delay exceeds 2.5s. As in Experiment 1, choices for distracter items C&D never exceeded 20% in any of the conditions; again well below the chance level of 50% (all $ps < 0.001$ on a Binomial test).

Statistical analyses corroborate these qualitative impressions. The proportion of choices for A was significantly higher than choices for B in the 2500ms condition, with 48 participants choosing A and only 15 choosing B (2 chose C) $Z = 4.03$, $p < 0.001$ on a Sign test; no significant difference in choice patterns was obtained in the 3250ms and 4000ms conditions, A: 34 participants, B: 27 participants, C: 4 participants, $Z = 0.77$ and A: 33 participants, B: 27 participants, C: 5 participants, $Z = 0.65$, respectively. Degradation of contiguity significantly affected choices for A, $Q(2,65) = 15.63$, and choices for B, $Q(2,65) = 9.60$, in the direction of decrease for A, and increase for B. Choices for the distracter items were not affected by variations in A's contiguity, $Q(2,65) = 2.00$, *n.s.*

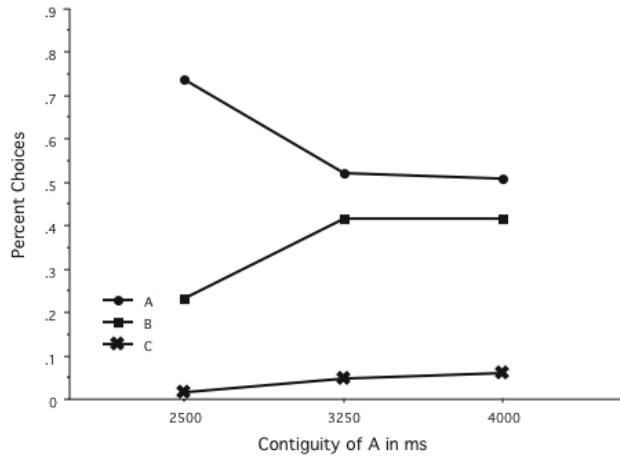


Figure 3. Experiment 2: Percentage of Participants (N=65) choosing Causes A or B or one of the two unrelated events (C & D).

As expected, with a steeper contiguity gradient, contingency no longer consistently dominated choice patterns. Both cues were equally important in determining choice patterns. Remarkably, B never was preferred over A, suggesting that contiguity never was more important than contingency, at least not within the parameters of this design.

Experiment 3

Until now, we have considered the low-contingency cause B associated with a 500ms delay as the contiguous alternative to the high-contingency cause A, which involved delays of 500ms and more. In many ways, 500ms could be deemed non-contiguous. We have argued elsewhere (Buehner & McGregor, 2006) that notions of contiguity are highly context dependent: a delay of 500ms might seem instantaneous when waiting for an elevator after having pressed the call button; that same delay might appear non-contiguous when flicking a lightswitch to illuminate a room. There are no clear guidelines as to what interval is interpreted as contiguous under what circumstances. The closest relevant evidence comes from Michotte's (1946/1963) studies on perceptual causality, where intervals of around 100-200ms were perceived as creating a temporal gap between cause and effect. Although our experiments are in a very different context from Michotte's, it is nonetheless worthwhile to check whether our general pattern of results – a general preference of contingency over contiguity – holds when contiguity is stronger. Experiment 3 thus increased the contiguity contrast between the high- and low-contingent causes even more by setting the contiguity of B to 0ms (instantaneous), while we employed the contiguity values of Experiment 1 for A.

Method

Participants

Sixty undergraduate students from Cardiff University participated to fulfill part of a course requirement.

Apparatus, Procedure, and Design

The same design and procedure as in Experiment 2 was employed, except that the contiguity of B was set to 0ms (instantaneous), while A could take on values of 500ms, 1000ms, 1500ms, and 2000ms. Trial lengths for these conditions were 4.17s, 4.67s, 5.17s, and 5.67s.

Results and Discussion

Table 4 lists the percentages of choices for the high (A) and low (B) probability cause, as well as the two non-causal distracter items (C) across the five probability gradients, and the four levels of contiguity for A. Inspection of Table 4 suggests that despite the increase in contiguity contrast between A and B, A still tends to be preferred over B, although the pattern tends to break down somewhat when the delay between A and the effect reaches 2s.

Table 4

Experiment 3: Percentages of choice for the high-probability cause (A), the low probability cause (B), or one of the two non-causal events (C), broken down by between-subject probability gradient (across) and within-subject contiguity-gradient for A (down). P_A and P_B refer to the probabilities that A or B is followed by the effect.

		P_A	0.75	0.75	1	1	1
		P_B	0.25	0.5	0.25	0.5	0.75
A	500ms		0.77	0.80	0.77	0.83	0.75
	1000ms		0.62	0.50	0.85	0.83	0.50
	1500ms		0.69	0.70	1.00	0.67	0.58
	2000ms		0.62	0.40	0.77	0.75	0.58
B	500ms		0.15	0.20	0.15	0.17	0.25
	1000ms		0.38	0.40	0.15	0.17	0.50
	1500ms		0.31	0.30	0.00	0.25	0.42
	2000ms		0.38	0.60	0.15	0.17	0.42
C	500ms		0.08	0.00	0.08	0.00	0.00
	1000ms		0.00	0.10	0.00	0.00	0.00
	1500ms		0.00	0.00	0.00	0.08	0.00
	2000ms		0.00	0.00	0.08	0.08	0.00

Figure 4 plots the percentages of choices for A, B, and C, collapsed over the probability gradients. A similar pattern to Experiment 1 emerges. Overall, A seems to be consistently preferred over B, but the extent of this preference diminishes, as contiguity for A decreases. Choices for the distracter items were well below 20% throughout, and thus below the chance level (all $ps < 0.001$ on a Binomial test). Statistical analyses confirm that A was chosen significantly more

often than B. For the contiguity of A of 500ms, 47 participants chose A, while 11 chose B (2 chose C), $Z = 4.60$; with 1000ms, 40 participants chose A, and 19 chose B (1 chose C), $Z = 2.60$; for 1500ms, 44 participants chose A and 15 chose B (1 chose C), $Z = 3.65$; finally, in the 2000ms condition, 38 participants chose A and 20 chose B (2 chose C), $Z = 2.23$. This last contrast, while still producing an individually significant result ($p = 0.026$), failed to produce a significant preference for A over B when Bonferroni alpha-correction ($p = 0.0125$) is applied.

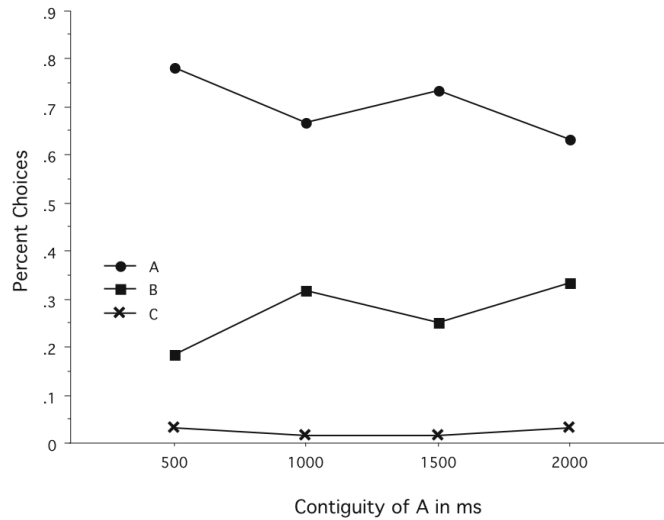


Figure 4. Experiment 3: Percentage of Participants (N=60) choosing Causes A or B or one of the two unrelated events (C & D).

Analyses for the overall impact of contiguity contrast on choices for A reveal that the degree of preference for A was not significantly affected by variations in A's contiguity, $Q(3,60) = 5.18$, *n.s.* Correspondingly, choices for B were likewise unaffected by variations of the contiguity contrast, $Q(3,60) = 5.59$, *n.s.* As in all previous experiments, participants' choices for the distracted items were unrelated to the contiguity contrast, $Q(3,60) = 0.67$, *n.s.* It seems as though our finding of a preference for a non-contiguous but high-contingent cause over a high-contiguous cause with poorer contingency is robust, and replicated even when contiguity for the high-contiguous cause was set to the maximum level possible (0ms).

Experiment 4

One potential criticism of our studies so far is that the experimental task might have put undue focus on contingency, at the expense of contiguity. More specifically, the forced choice selection of one UFO to open the gate might have rendered variations in contiguity inconsequential: by asking participants to choose one candidate in order to produce the effect, we might have implicitly shifted the

focus on contingency, and made variations in contiguity appear meaningless for the task. We address this issue in Experiments 4 and 5. In Experiment 4, we examine new probability combinations. We were particularly interested to obtain a “manipulation check” to make sure our paradigm can elicit choices based on a preference for the more contiguous cause even in the forced choice task. To this end, we created four new probability contrasts, three of which employed identical probabilities of success for UFOs A and B to open the gate, while preserving the structure implemented in Experiments 1-3, namely that B is always the contiguous cause. If contiguity were inconsequential to our task, then variations in the contiguity contrast, while keeping the probabilistic contrast constant should have no systematic impact on choice patterns; participants should be impartial between A and B. On the other hand, if our task is sensitive to variations in contiguity, then choices should favour B over A in these conditions. A fourth contrast employed the probabilities 0.5 and 0.25 for A and B, respectively. This was included to investigate the contingency advantage in low probability conditions. More specifically, the “winning” contingency in Experiment 1-3 was always very high (1 or 0.75). Perhaps such high contingencies, especially when pitted against lower ones (0.25 or 0.5), were especially likely to yield preference for contingency over contiguity.

Method

Participants

125 undergraduate students from Cardiff University participated to fulfill part of a course requirement.

Apparatus, Procedure, and Design

The same procedure and design as in Experiment 1 was employed, except that the following four new probability values for A and B were employed: a) 50% vs. 50%, b) 75% vs. 75%, c) 100% vs. 100%, d) 50% vs. 25%. The same contiguity gradients as in Experiment 1 were employed, i.e. the delay associated with B was always set to 500ms, while the delay for A could take on values of 500ms, 1000ms, 1500ms, and 2000ms. The event configurations were altered as necessary to maintain true independence of A and B's respective probabilistic relation to the effect. Participants were randomly assigned to one of the four probability gradients.

Results and Discussion

Table 5 lists the percentages of choices for A and B, as well as the two non-causal distracter items (C) across the four probability gradients, and the four levels of contiguity for A. Inspection of Table 5 suggests that when the contingency contrast favours neither of the two candidates, participants base their choice on contiguity: a clear preference of B over A emerges, when the contiguity contrast shifts against A. The results also suggest that a contingency advantage of A over B (despite a contiguity contrast favoring B) seems to hold even when the contingencies involved are small and the contrast itself marginal (0.5 vs 0.25).

Table 5

Experiment 4: Percentages of choice for A, B, or one of the two non-causal events (C), broken down by between-subject probability gradient (across) and within-subject contiguity-gradient for A (down). P_A and P_B refer to the probabilities that A or B is followed by the effect.

		P_A	0.50	0.75	1.00	0.50
		P_B	0.50	0.75	1.00	0.25
A	500ms		0.30	0.44	0.48	0.81
	1000ms		0.27	0.47	0.26	0.63
	1500ms		0.37	0.41	0.13	0.56
	2000ms		0.30	0.25	0.19	0.59
B	500ms		0.60	0.53	0.39	0.06
	1000ms		0.67	0.47	0.61	0.28
	1500ms		0.47	0.59	0.84	0.34
	2000ms		0.53	0.75	0.74	0.34
C	500ms		0.10	0.03	0.13	0.13
	1000ms		0.07	0.06	0.13	0.09
	1500ms		0.17	0.00	0.03	0.09
	2000ms		0.17	0.00	0.06	0.06

Figure 5 plots the percentages of choices for A, B, and C, collapsed over the three probability gradients where $P_A = P_B$. As in the previous experiments, choices for C were below 20% throughout, and thus well below the chance level (all $ps < 0.001$). As expected, with identical contiguities of 500ms, choices do not reliably discriminate between A and B, when both have the same probability to produce the effect: 47 participants chose A, while 38 chose B (8 chose C), $Z = 0.868$, *n.s.* As soon as the contiguity contrast favors B, however, choices reflect this accordingly. With contiguity of A of 1000ms, 31 participants chose A, and 54 chose B (8 chose C), $Z = 02.39$, $p = 0.017$, with contiguity of A of 1500ms, 28 participants chose A, and 59 chose B (6 chose C), $Z = 3.22$, $p < 0.001$, and with contiguity of A of 2000ms, 23 participants chose A, and 63 chose B (7 chose C), $Z = 4.21$, $p < 0.001$. The preference of B over A in the 500ms vs 1000ms contiguity contrast is significant when considered alone, but not with Bonferroni alpha correction ($p = 0.0125$). Cochran's Q tests revealed that the preference of B over A does not vary as a function of the extent of the contiguity contrast in favor of B: The increase in choices for B falls short of significance, $Q(3) = 6.77$, $p = 0.08$. Correspondingly, A choices do not seem to decrease, $Q(3) = 5.90$, *n.s.*; as expected, choices for C are unaffected by variations in contiguity contrast, $Q(3) = 0.62$, *n.s.*

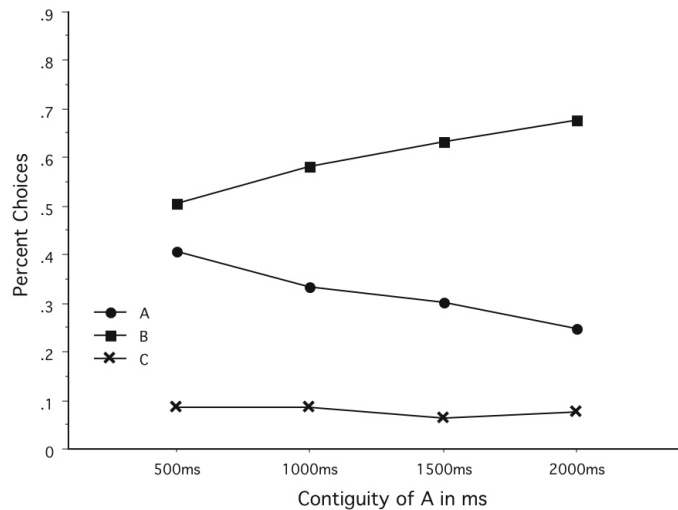


Figure 5. Experiment 4: Percentage of Participants (N=93) choosing Causes A or B or one of the two unrelated events (C&D) when Probabilities associated with A and B were identical.

Analysis of the data from the 32 participants who were in the 50% vs 25% probability gradient condition reveals that, as expected, the higher probability of A leads to a preference of A over B when contiguity is identical at 500ms in both: 26 participants chose A, and only 2 chose B (4 chose C), $Z = 04.35$, $p < 0.001$. This clear preference shows that the sensitivity to contingency is not limited to high probabilities, or a function of clear contingency gradients. People seem to be able to detect and exploit small differences in contingencies, even when the contingencies involved are quite low. The contingency advantage of A over B fades, though, when B gains a contiguity advantage. When A's contiguity was 1000ms, 20 participants chose A, and 9 chose B (3 chose C), $Z = 1.86$. While this result would be significant with $p = 0.063$ as an individual test, it falls short of the more conservative Bonferroni adjusted value of 0.0125. With A's contiguity degrading further to 1500 and 2000ms, the pattern weakens further: 18 participants chose A, 11 chose B (C: 3) for the former, and 19 for A, 11 for B (C: 2) for the latter. Note that this pattern still reflects a notional preference for A over B, and that the pattern never reverses towards a preference of B over A. Cochran's Q tests revealed that the trend from A choices towards B, although nominally there, falls short of significance: For the B choices, $Q(3) = 8.32$, $p = 0.04$ (not significant after Bonferroni correction); correspondingly, for A choices, $Q(3) = 5.47$, *ns*; as in previous experiments, there was no increase in C choices, $Q(3) = 1.09$, *ns*.

Experiment 5

Experiment 5 replicated Experiment 1, but employed a standard rating scale as the dependent variable. As mentioned before, forcing participants to choose the one UFO they would activate to open the gate might put undue

emphasis on contingency at the expense of losing any influence that contiguity might have had on subjective impressions of causality. Consequently, we substituted the forced choice method with four separate numerical ratings, one for each candidate cause. This change in method should facilitate any contiguity advantage which our previous experiments might have overlooked to shine through.

Method

Participants

Forty-six undergraduate students from Cardiff University participated to fulfill part of a course requirement.

Apparatus, Procedure, and Design

The same procedure and design as in Experiment 1 were employed, except that participants had to rate each UFO at the end of each condition, rather than pick just one. To this end, the following rating scale, was employed once a condition had finished: "Please rate each individual UFO in how successful it was in opening the gate. Provide a number between 0 and 100 where 0 means a UFO never succeeded in opening the gate, 50 means a UFO sometimes succeeded in opening the gate, and 100 means a UFO always succeeded in opening the gate". Participants were presented with this scale on an otherwise blank screen, and were prompted to press RETURN to rate each UFO individually. After this, the same screen as during the trial (see Figure 1) was presented, with a text box underneath each UFO, and additional instructions to click on the Stargate at the center once each UFO had been rated. Participants entered individual ratings in each corresponding box, and clicked the gate when they had completed the rating to proceed. The same probability gradients as in Experiment 1 were varied between subjects: 75% vs. 25%; 75% vs. 50%; 100% vs. 25%; 100% vs. 50%; and 100% vs. 75%. Contiguity for B was always at 500ms, while A's contiguity could take on values of 500ms, 1000ms, 1500ms, and 2000ms. Variations of A's contiguity were implemented within subjects, resulting in each participant working on four conditions.

Results and Discussion

Each participant provided four ratings per condition, one for option A (high probability – low contiguity), one for B (low probability – high contiguity), and one for each of the two unrelated distracter items. These latter two ratings were collapsed into one score for each participant and condition by simply taking the average of the two. Data for the 500ms condition were lost for one participant in the 75% vs 25 % probability gradient group due to a computer error and were treated as missing in the analyses.

Figure 6 plots the mean causal ratings for the two competing causal candidates and the distracter items as a function of A's contiguity, and broken down by contingency gradient. Inspection of Figure 6 suggests that option A – the one with the higher probability – is consistently rated as more causal than option B – the one with better contiguity, and that this preference remains stable even when the contiguity advantage of B becomes more pronounced. The exception to this pattern is found in Panel b, which shows results for the 75% vs 50% gradient; here A and B are equally rated when A's delay is larger than 1000ms.

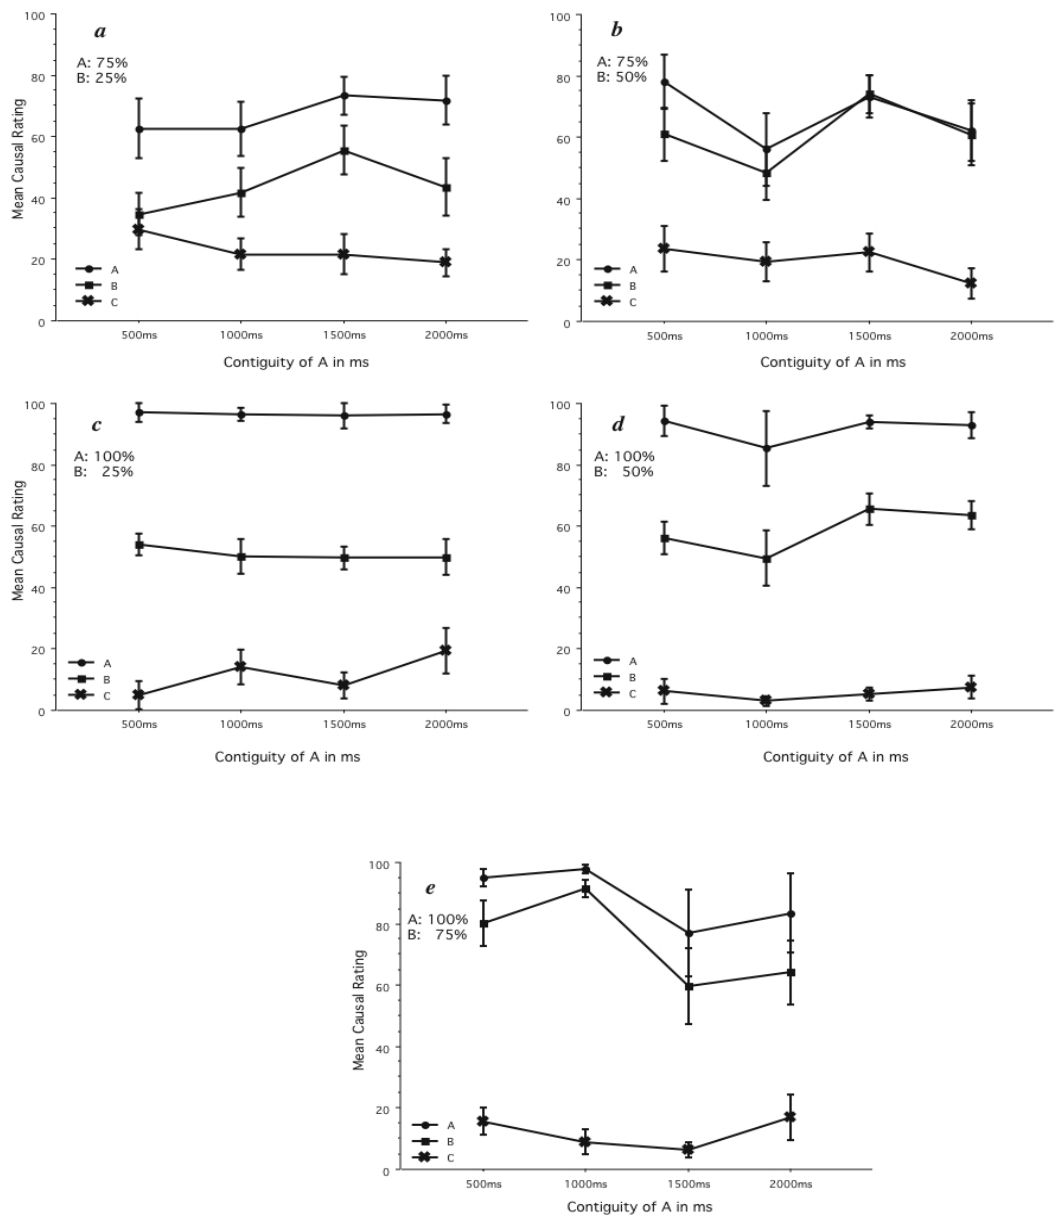


Figure 6. Experiment 5: Mean Causal Ratings for Causes A, B, and the two unrelated events (C&D), broken down by contiguity gradient. a: 75% vs 25% (N=10), b: 75% vs 50% (N=10), c: 100% vs 25% (N=10), d: 100% vs 50% (N=8), e: 100% vs 25% (N=7). Error bars indicate standard error.

Statistical analyses corroborate these observations. A 3 (Candidate Cause) x 4 (Contiguity for A) x 5 (Contiguity Gradient) mixed ANOVA revealed a main effect of *Candidate Cause*, $F(2,80) = 289.20$, $MSE = 56804.74$, $\eta^2 = 0.878$, and a

Contingency Gradient x Candidate Cause Interaction, $F(8,80) = 8.152$, $MSE = 46309.00$, $\eta^2 = 0.449$, reflecting that ratings for A were higher compared to B, and that both varied as a function of their respective contingency. The only significant effect associated with *Contiguity for A* was the *Contiguity for A x Contingency Gradient* interaction, $F(12,120) = 2.13$, $MSE = 10808.21$, $\eta^2 = 0.177$.

These results suggest that the contingency advantage reported in the previous experiments is a robust finding, and not an artifact of the forced choice method. It appears that people put more emphasis on contingency than on contiguity when they evaluate the causal efficacy of cues in a knowledge lean paradigm.

General Discussion

The goal of this paper was to investigate how contiguity and contingency relate to each other in causal induction. In particular, we wanted to find out whether people selectively weigh one cue as more important than the other in situations where no prior knowledge is available to bias one timeframe over another. Towards this end we created a novel experimental setup, which allowed us to study causal induction under ecologically valid conditions: events were presented in one continuous flow, with no discrete trial boundaries. This allowed us to focus on our main question of interest: would people put equal emphasis on cause-effect contiguity and contingency? Or would one cue receive more weight in determining the best causal candidate?

Previous experiments investigating the role of contiguity within demarcated learning trials examined contrasts between experienced and expected timeframes (e.g., Allan et al., 2003). In such studies, variations of contiguity determined whether each individual *trial* was seen as evidence for or against the causal relation in question. In our design, the absence of trials made such evaluations immaterial. Instead, participants had to consider the entire stream of events when making causal judgments, and had to decide which of four potential causes was the strongest. In making these choices, participants had to consider both cause-effect contingency and contiguity.

When considering such trade-offs, it is important to separate *utility* from *causality* (Oaksford & Chater, 1998). A response may have perfect contingency with a desired outcome, but produce the outcome only after a long delay. An alternative response may produce the outcome right away, but unreliably. Depending on the cost of responding and the time available to interact with the environment, it may be more beneficial for the organism to engage in the low-contingent, but high-contiguity response. Experiments 1 through 4 asked participants to select the cause that is most successful in producing the effect. The one-shot nature of this dependent measure clearly requested an answer based on causality, rather than utility; our results show not only that participants were aware of this, but that – by and large – their assessments of causality were dominated by the cause-effect contingency. .

One potential criticism of this paradigm is that its task demands put more emphasis on probability than on contiguity. For example, had we instructed participants that an enemy UFO is approaching and they need to open the gate as quickly as possible to escape, then participants might have chosen the high contiguity cause more often. We would certainly expect that task demands can raise the subjective importance of contiguity or contingency. But nothing in our instructions actually highlighted contingency as more important than contiguity. We certainly did not instruct people that they should choose the UFO that opens the gate *irrespective* of time, or that achieving the outcome at all is more important than when it occurs. So if our results indicate that contingency receives more weight than contiguity in a knowledge-lean paradigm, this probably reflects that contingency by default is more important than contiguity.

We initially deliberately avoided traditional causal ratings in order to make it clear to our participants that we wanted to know which of the four alternatives they thought was the strongest cause. We were concerned that any dominance of one cue over the other might be so subtle as to be swamped if we allow participants to evaluate each cause individually via a rating. More specifically, we thought that participants might trade-off deficits on one dimension against advantages on the other, and that numerical ratings would thus not reflect any dominance patterns. However, Experiment 5 demonstrated that the contingency advantage was more pervasive than that. Even when each cause could be individually appraised so that each rating could be a function of contingency and contiguity, the contingency advantage prevailed. No trade-off between the two cues was found at all, within the 2s span that we considered – again it appeared that contingency dominated contiguity.

Despite the overall finding that contingency dominated contiguity in our experimental paradigm, several caveats are in order. Firstly, Experiments 1 and 2 revealed that the contingency advantage appears to fade away if the contiguity contrast is large enough. More specifically, a high probability cause associated with a delay of more than 2.5s was no longer seen as reliably more attractive than a low probability alternative associated with a short delay. While this may be indicative of a genuine benefit of contiguity, it may ironically also be even stronger evidence for the domination of contingency over contiguity: As discussed in the introduction, one consequence of long cause-effect delays is that it increases the probability of intervening alternative causes, which then might compete for explanatory strength with the candidate cause. Our experimental paradigm, which ensured complete independence of all candidate causes, deliberately allowed this possibility. Consequently, in conditions where the delay associated with A was long, the probability of B (or any of the two distracter items) to occur between A and its associated effect would have been increased, resulting in event sequences of the nature A---B---E, where E represents the effect. Unless participants had learnt the delay associated with A already, there would be no way of knowing that the effect was due to A, and they might have (erroneously) attributed it to B instead. Thus, the apparent benefits of cause-effect contiguity may not actually reflect an advantage of B's high contiguity, but instead a subjective boost of B's

contingency, resulting from A's degraded contiguity. However, given that – at least within the boundaries of our paradigm – we have never found a domination of contiguity over contingency, this does not detract from our main finding: that contingency appears to be more fundamental to causal learning than contiguity.

We hope that the empirical results presented in this paper will inform and constrain modeling efforts in causal learning in real time, and are seen as the beginning of a research program towards causal reasoning and learning in continuous time. Our current results suggest that humans might be inherently biased to weigh probability stronger than contiguity. Naturally, this hierarchy must be contingent on reasoners still being able to *notice* the contingency. If cause-effect delays exceed the perceptual or cognitive threshold of the organism (which undoubtedly will be task-dependent), then contiguous but low-contingent alternatives will subjectively appear to be the only candidates.

References

- Allan, L. G., & Jenkins, H. M. (1980). The judgment of contingency and the nature of response alternatives. *Canadian Journal of Psychology*, *34*(1), 1-11.
- Allan, L. G., Tangen, J. M., Wood, R., & Shah, T. (2003). Temporal contiguity and contingency judgments: A Pavlovian analogue. *Integrative Physiological and Behavior Science*, *31*(2), 205-211.
- Anderson, J. R., & Sheu, C. F. (1995). Causal inferences as perceptual judgments. *Memory and Cognition*, *23*(4), 510-524.
- Buehner, M. J., & Cheng, P. W. (2005). Causal Learning. In K. J. Holyoak & R. Morrison (Eds.), *Handbook of thinking & reasoning* (pp. 143-168). Cambridge, England: Cambridge University Press.
- Buehner, M. J., & May, J. (2002). Knowledge mediates the timeframe of covariation assessment in human causal induction. *Thinking and Reasoning*, *8*(4), 269-295.
- Buehner, M. J., & May, J. (2003). Rethinking temporal contiguity and the judgment of causality: Effects of prior knowledge, experience, and reinforcement procedure. *Quarterly Journal of Experimental Psychology*, *56A*(5), 865-890.
- Buehner, M. J., & McGregor, S. (2006). Temporal delays can facilitate causal attribution: Towards a general timeframe bias in causal induction. *Thinking and Reasoning*, *12*(4), 353-378.
- Cheng, P. W. (1997). From covariation to causation: A causal power theory. *Psychological Review*, *104*(2), 367-405.
- Einhorn, H. J., & Hogarth, R. M. (1986). Judging probable cause. *Psychological Bulletin*, *99*(1), 3-19.
- Greville, W. J., & Buehner, M. J. (2007). The influence of temporal distributions on causal induction from tabular data. *Memory & Cognition*, *35*(3), 444-453.
- Griffiths, T. L., & Tenenbaum, J. B. (2005). Structure and strength in causal induction. *Cognitive Psychology*, *51*(4), 285-386.
- Holyoak, K. J., & Hummel, J. E. (2000). The proper treatment of symbols in a connectionist architecture. In E. Dietrich & A. Markman (Eds.), *Cognitive dynamics: Conceptual change in humans and machines*. (pp. 229-263). Mahwah, NJ: Erlbaum.
- Hume, D. (1739/1888). A treatise of human nature. In L. A. Selby-Bigge (Ed.), *Hume's treatise of human nature*. Oxford, UK: Clarendon Press.

- Jenkins, H., & Ward, W. (1965). Judgment of contingencies between responses and outcomes. *Psychological Monographs*, 7, 1-17.
- Mandel, D. R., & Lehman, D. R. (1998). Integration of contingency information in judgments of cause, covariation, and probability. *Journal of Experimental Psychology: General*, 127(3), 269-285.
- Mendelson, R., & Shultz, T. R. (1976). Covariation and temporal contiguity as principles of causal inference in young children. *Journal of Experimental Child Psychology*, 22(3), 408-412.
- Michotte, A. E. (1946/1963). *The perception of causality* (T. R. Miles, Trans.). London: Methuen & Co.
- Novick, L. R., & Cheng, P. W. (2004). Assessing interactive causal influence. *Psychological Review*, 111(2), 455-485.
- Oaksford, M., & Chater, N. (1998). *Rational models of cognition*. New York: Oxford University Press.
- Pearl, J. (2000). *Causality: models, reasoning, and inference*. Cambridge, England: Cambridge University Press.
- Reichenbach, H. (1956). *The direction of time*. Berkeley & Los Angeles: University of California Press.
- Rescorla, R. A., & Wagner, A. R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In A. H. Black & W. F. Prokasy (Eds.), *Classical conditioning II: Current theory and research* (pp. 64-99). New York: Appleton-Century Crofts.
- Savastano, H. I., & Miller, R. R. (1998). Time as content in Pavlovian conditioning. *Behavioural Processes*, 44(2), 147-162.
- Schlottmann, A. (1999). Seeing it happen and knowing how it works: How children understand the relation between perceptual causality and underlying mechanism. *Developmental Psychology*, 35(5), 303-317.
- Shanks, D. R., & Dickinson, A. (1987). Associative accounts of causality judgment. In G. H. Bower (Ed.), *Psychology of learning and motivation-advances in research and theory* (Vol. 21, pp. 229-261). San Diego, CA: Academic Press.
- Shanks, D. R., Pearson, S. M., & Dickinson, A. (1989). Temporal contiguity and the judgment of causality by human subjects. *Quarterly Journal of Experimental Psychology Section B- Comparative and Physiological Psychology*, 41(2), 139-159.
- Shultz, T. R. (1982). Rules of causal attribution. *Monographs of the Society for Research in Child Development*, 47(1), 1-51.
- Siegler, R. S., & Liebert, R. M. (1974). Effects of contiguity, regularity, and age on children's causal inferences. *Developmental Psychology*, 10(4), 574 -579.
- Steyvers, M., Tenenbaum, J. B., Wagenmakers, E. J., & Blum, B. (2003). Inferring causal networks from observations and interventions. *Cognitive Science*, 27(3), 453-489.
- Tenenbaum, J. B., & Griffiths, T. L. (2001). Structure learning in human causal induction. In T. G. Dietterich & V. Tresp (Eds.), *Advances in neural processing systems* (Vol. 13, pp. 59-65). Cambridge, MA: MIT Press.
- Wasserman, E. A., & Neunaber, D. J. (1986). College students' responding to and rating of contingency relations: The role of temporal contiguity. *Journal of the Experimental Analysis of Behavior*, 46(1), 15-35.
- White, P. A. (2003). Making causal judgments from the proportion of confirming instances: The pCI rule. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 29(4), 710-727.

- Young, M. E. (1995). On the origin of personal causal theories. *Psychonomic Bulletin and Review*, 2(1), 83-104.
- Young, M. E. (1997). Implicit processes in the development of causal knowledge: A connectionist model of the use of Humean cues. In P. W. van den Broek, P. J. Bauer & T. Bourg (Eds.), *Developmental spans in event comprehension and representation: Bridging fictional and actual events*. (pp. 29-50). Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.