(in press). In T. Shipley & J. Z In T. Shipley & J. Zacks (Eds.), *Understanding events: How humans see, represent, and act on events*. Oxford University Press.

Dynamics and the perception of causal events

Phillip Wolff

Emory University

Please address correspondence to:

Phillip Wolff

Department of Psychology

Emory University

Department of Psychology

532 N. Kilgo Cir.

Atlanta, GA 30322

Office #: 404-727-7140

Fax #: 404-727-0372

Dynamics and the perception of causal events

To imagine possible events, we use our knowledge of causal relationships. To look deep into the past and infer events that were not witnessed, we use causal relationships as well. We also use causal knowledge to infer what can not be directly seen in the present, for instance, the existence of planets around distant stars, or the presence of subatomic particles. Knowledge of causal relationships allows us to go beyond the immediate here and now. In this chapter I introduce a new theoretical framework for how this very basic concept might be mentally represented.

In effect, I propose an epistemological theory of causation, that is, a theory that specifies the nature of people's knowledge of causation, the notion of causation used in everyday language and reasoning. In philosophy, epistemological theories are often contrasted with metaphysical theories, theories about the nature of reality. Since people's concepts of causation are assumed to be in error, most metaphysical theories of causation seek to reform rather than describe the concept of CAUSE in people's heads, (see Mackie, 1974; Dowe, 2000). Theories of causation in psychology have followed suit by linking people's representations of causation to the outward manifestations of causation rather than to the quantities in the world that produce those manifestations. In this chapter, I explore another possibility, one that assumes that people can peer beyond the veil of the visible to represent the (invisible) elements of the world that are essential to causal events. In the theory of causation to be presented, I propose that while the essential elements of causation are invisible, they are also are highly inferable because they are lawfully connected to the visible properties of events. As a consequence, people's representations of causation partially reproduce causation in the world. In short, I propose an epistemological theory that could, in many respects, also serve as a metaphysical theory.

Theories of Causation of the Humean Tradition

The distinction between causation in the mind and in the world is discussed in Hume (1737/1975). He noted that when people first observe a causal relationship, they may be able to detect spatial-temporal contiguity, a succession of events and, in particular, covariation, but not what is most central to people's ordinary concept of causation, that is, force, necessary connection, causal power and/or energy. Since notions such as these cannot be determined from observation, let alone from reasoning, Hume held that people's ordinary notion of causation had no clear basis in reality (Mackie, 1974; Dowe, 2000).

Hume's arguments have greatly influenced psychological theories of causation. The invisibility of causation has led many researchers to conclude that causal relationships can only be induced from event frequencies. Induction from event frequencies implies that causal relations are represented in terms of their observable outcomes, rather than in terms of the physical quantities that actually produce those outcomes. Such theories are exemplified by, for example, Cheng and Novick's (1991, 1992) *probabilistic contrast model*. In this theory, causal relations are based on the covariation people observe between a candidate cause and effect within a "focal set" of events. Specifically, facilitative causation (CAUSE) is inferred when the probability of the effect in the presence of a candidate cause, $P(E|C)$, is noticeably greater than the probability of the effect in the absence of the cause, $P(E|\neg C)$, that is, when the difference between these two probabilities, ΔP, is positive. Inhibitory causation (PREVENT) is inferred when ΔP is negative. While the probabilistic contrast model is primarily a model of causal induction, it also implies a theory of causal representation in which causal relationships are associated with statistical dependencies. Cheng's (1997) power PC theory of causation extends the probabilistic contrast model by proposing that people's causal judgments are based on a theoretical entity—causal power—that can be estimated from covariation, provided certain conditions are honored (see Luhmann & Ahn, 2005 for a critical analysis of these assumptions). The power PC model fleshes out Hume's (1748/1975) intuition that people's ordinary notions of causation involve the notion of causal power and that this notion is derived from covariational information. Importantly, however, causal power in the power PC model is still very much a statistical entity: it is determined purely on the basis of co-occurrence patterns across multiple events, not on the physical quantities in the world that bring about those patterns (just as Hume would have liked).

Another recent approach to representation of causal relationships is captured in Bayesian network models of causation. In Bayesian models, causal factors are linked together in a network of nodes and arrows indicating causal connections (Gopnik, Glymour, Sobel, Shultz, Kushnir, & Danks, 2004; Pearl, 2000; Sobel, Tenenbaum, & Gopnik, 2004; Sloman & Lagnado, 2002; Sloman, 2005; Tenenbaum & Griffiths, 2001). While Bayesian networks go beyond prior approaches to causation in being able to address phenomena associated with causal reasoning, they remain very much theories based only on the visible. Before a Bayesian network can reason, it must first be constructed, and all current approaches to the construction of Bayesian networks involve determining statistical dependencies (Wolff, 2006). As stated above, statistical dependencies are summary representations that are tabulated and stored in people's minds, not in the world. Such networks do not directly represent the processes and quantities that bring about causal relationships in the world. As pointed out by Bunge (1959), to treat statistical dependency approaches as theories of causation is to confuse causation for one of it tests.

In addition to probability approaches to causation, Hume's views have influenced accounts of causal perception, most notably Michotte's (1946/1963). Michotte examined in great detail the stimulus conditions that give rise to the impression of causation. He focused primarily on what he called launching events. In the canonical launching event, an object A approaches and hits a stationary object B, sending it into motion. One of the main findings from this research is that even small changes in the stimulus conditions could greatly affect the impression of causation (Scholl & Tremoulet, 2000). For example, a temporal gap between the two objects can sometimes disrupt the impression of causation (Michotte, 1946/1963; Thinès, Costall, & Butterworth, 1991). Importantly, when people are asked to describe launching events, their descriptions typically include more than a mere specification of the objects' motions (Choi & Scholl, 2005; Michotte, 1946/1963). In billiard-ball collisions (a type of launching event), for example, people not only see changes in motion in two objects but also view one of the balls as causing the other's motion. The pattern of motion instantiated by the launching event leads to the apprehension of an invisible causal agency. However, Michotte (1946/1963) emphasized that this causal agency was not a direct representation of forces or energies in the world. Rather, it was a perceptual phenomenon, or illusion, possibly triggered by an innate perceptual mechanism.

Challenges for theories of the Humean tradition

Many current models of causation are based on Hume's assumption that the physical quantities that give rise to causal events are unavailable to the human perceiver. But on this assumption, I will argue, Hume was mistaken. While people cannot directly see forces and energies, this does not mean that such quantities cannot be directly recovered from the sensory input. Newtonian physics indicates that, in principle, such recovery is possible because of the lawful mapping between kinematics and dynamics. Kinematics specifies the observable properties of an event: the shapes, sizes, spatial relations, velocities, and accelerations of the various entities in a motion event (Schwartz, 1999; Joskowicz & Sacks, 1991; Gilden, 1991). In contrast, a motion event's dynamics concerns the invisible properties of an event (Schwartz, 1999), specifically, the forces, potential energies, pressures, powers, elasticities, and masses that bring about an event. Some of the mapping between kinematics and dynamics is captured in Newton's laws of motion. For example, if an object suddenly turns to the right, Newton's 1st law states that the change in velocity implies acceleration, which entails the presence of a force. Newton's 2nd law, $\mathbf{F} = m\mathbf{a}$, implies that the direction of the force, $\mathbf{F}$, is exactly the same as the direction of acceleration, $\mathbf{a}$. Thus, by observing an instance of change in velocity and the direction of that change, people can, in principle, detect the presence of a force and the direction of its influence. The process of computing forces from kinematics is known as *inverse dynamics.*

Within the field of physics understanding, there is general agreement that people are capable of performing inverse dynamics, at least to a limited extent (Hecht, 1996; Gilden 1991; Kaiser & Proffitt, 1984; Kaiser, Proffitt, Whelan, & Hecht, 1992; Proffitt & Gilden, 1989; Twardy & Bingham, 2002). According to Runeson and his colleagues, people's ability to infer the dynamic properties of an event is quite good, because people's perceptual systems allow them to "see" the dynamics of an event via its kinematics, a proposal known as the principle of kinematic specification of dynamics (KSD; see Runeson & Frykolm, 1983; Runeson, Juslin, & Olsson, 2000; Runeson & Vedeler, 1993). Others have been more conservative in their estimates of people's ability to infer the dynamics of an event, suggesting that people may be able to recover a portion of these properties via perceptual heuristics (see Proffitt & Gilden, 1989; Hecht, 1996; Gilden, 1989). Clearly, people's ability to infer dynamic properties is not perfect (e.g., McCloskey, 1983; McCloskey & Kohl, 1983; McCloskey, Washburn, & Felch, 1983). People sometimes fail to notice certain kinds of dynamic properties or impute properties that do not exist (Clement, 1983; McCloskey, 1983). Nevertheless, the process of inducing dynamic properties is not completely arbitrary. This is especially true when the dynamic properties of a situation do not depend on the geometry of the objects, in which case the moving entities in a scene can be treated as particles (Gilden, 1991; Proffitt & Gilden, 1989; Hecht, 1996; Kaiser, Proffftt, Whelan, & Hecht, 1992).

Beyond the empirical literature, our ability to infer the presence of dynamic properties is supported by everyday experience. Human bodies can detect energy. When we touch a hot pan, we not only feel the solidness of the pan, but also the heat it gives off. Bigelow, Ellis, and Pargetter (1988) provide a similar example in the case of forces. If something bumps us and we stumble, we feel the force. We can argue that it is the force that is felt—not only the object—because the same object feels different when it bumps us gently or hard. Hume's assumption that people do not have access to the agencies that bring about causal events is inconsistent with work in physics understanding as well as with common sense. Just because forces and energies cannot be seen does not mean that their existence cannot be sensed in other ways, or that their presence cannot be inferred, relatively directly, from visual input. Therefore, one of the basic assumptions motivating models in the Humean tradition is false. People have at least partial access to the quantities in the world that bring about causation in the world.

Physicalist models of causation

In contrast to Humean models, physicalist models hold that people's representations of causation may partially copy or reproduce what goes on in the real world. The basic idea in physicalist approaches to causation is that people's representations of causation specify physical quantities in the world, such as energy, momentum, linear and angular momentum, impact forces, chemical forces, and electrical forces, among others. For example, according to Aronson's (1971) Transference Theory, causation implies contact between two objects in which a quantity possessed by the cause (e.g., velocity, momentum, kinetic energy, heat, etc.) is transferred to the effect. Another transference theory is proposed by Fair (1979), who holds that causes are the source of physical quantities, energy, and momentum that flow from the cause to the effect. According to Salmon's (1994, 1998) Invariant Quantity theory, causation involves an intersection of world lines that results in transmission of an invariant quantity. The proposals of Aronson, Fair, and Salmon come from the philosophy literature. Similar proposals from the psychology literature have been termed generative theories of causation. According to Bullock, Gelman, and Baillargeon (1982), adults believe that causes bring about their effects by a transfer of causal impetus. Shultz (1982) suggests that causation is understood as a transmission between materials or events that results in an effect. According to Leslie (1994), physical causation is processed by a "Theory of Bodies" (ToBy) that schematizes objects as bearers, transmitters, and recipients of a primitive notion of force.

A recent proposal from the philosophy literature breaks from earlier physicalist models in not requiring a one-way transmission of energy or momentum. According to Dowe's Conserved Quantity Theory (2000), there are two main types of causation: persistence (e.g., inertia causing a spacecraft to move through space) and interactions (e.g., the collision of billiard balls causing each ball to change direction). Causal interactions occur when the trajectories of two objects intersect and there is an *exchange* of conserved quantities (e.g., an exchange of momentum when two billiard balls collide). Unlike transfers, exchanges are not limited to a single direction (e.g., from cause to effect).

Assumptions of physicalist theories. Physicalist approaches to causation share several assumptions. First, they assume that an interaction can be identified as causal on the basis of properties that belong solely to that interaction. Second, defining causal relationships in terms of physical quantities imposes a relatively 'local' level of granularity on the analysis of causal relationships. Transfer of energy, for example, can only occur through local interactions between objects. Third, at the 'local' level of granularity, causal relationships are deterministic (Goldvarg & Johnson-Laird, 2001; Luhmann & Ahn, 2005): the physical quantities that instantiate direct causal relationships are either present or absent, not present to a probabilistic degree. Fourth, the 'local' nature of causal connections implies that when there is a causal connection between two non-contiguous events, there must be a causal chain of intermediate links, each contiguous to the next (Russell, 1948). Hence, physicalist theories imply the need for causal mechanisms, as has been supported by work in psychology (Ahn & Bailenson, 1996; Ahn & Kalish 2000; Ahn, Kalish, Medin, & Gelman, 1995; see also Bullock et al., 1982; Shultz 1982). The fifth commonality is that most physicalist theories reduce causal relationships to quantities that cannot be directly observed. In the language of physics, physicalist models hold that people represent causal relationships in terms of their dynamics rather than kinematics. The sixth assumption is that physical causation is cognitively more basic than non-physical causation (e.g., social or psychological causation). In support of this assumption, the ability to perceive physical causation begins to develop earlier in infants (around 3 to 4 months) than the ability to perceive social causation (around 6 to 8 months; Leslie, 1994; Cohen, Amsel, Redford, & Casasola, 1998; Oakes, 1994). A final assumption is that non-physical causation is in some way modeled after physical causation (Leslie, 1994; Talmy, 1988). This modeling may occur via a process of analogy in which notions such as "effort" and "intention" are construed as energies and forces.

Evaluation of physicalist accounts of causations. Physicalist models hold that people's representations of causation refer to physical quantities in the world. As a consequence, such theories can provide a precise characterization of the physical agencies that bring about causal events. They can also provide an account of how causal relationships might be inferred on the basis of a single observation. This is possible because the information needed to infer causal relationships is available in the occurrence of individual events. While physicalist models have several strengths, they also have several limitations. Arguably the most important test of a theory of causation is whether it has extensional adequacy. A theory of causation should be able to pick out the range of situations that people judge to be causal while excluding situations that people judge to be non-causal. However, current physicalist models do not categorize causal situations in the same way as people. In particular, the physicalist models discussed so far conflate the concepts of CAUSE and ENABLE. For people, these concepts are similar but not synonymous. In most contexts, they are not interchangeable, as illustrated by the sentences in (1) and (2).

(1) a. Hinges enabled the crutches to fold in half.
 b. A weak spot in the welding caused the crutches to fold in half.

(2) a. ?Hinges caused the crutches to fold in half.
 b. ?A weak spot in welding enabled the crutches to fold in half.

The sentences in (1a) and (1b) are perfectly acceptable. However, if the verbs in (1a) and (1b) are switched, the resulting sentences (2) sound odd. Importantly, however, in all of scenarios described in (1) and (2) there is a transmission or exchange of energy. Another limitation of current physicalist theories is that they do not easily represent the concept of PREVENT (Dowe, 2000). If prevention is characterized by the lack of transfer or exchange of energy, then it does not differ from the absence of any kind of interaction and if it is characterized by a transfer or exchange of energy, it does not differ from causation. (See Dowe (2000) for an in-depth discussion of the problem of prevention.) The problem with the physicalist models discussed so far is that transmission or exchange of energy is too coarse a criterion for distinguishing causation from other kinds of events that also involve a transmission or exchange of energy. To distinguish causation from other kinds of relationships, a finer level of representation is required.

The dynamics model

The *dynamics model* is a physicalist model of causation. As such, it holds that people represent causal relations in a manner that copies or reproduces the way in which causal relationships are instantiated in the real world. While it is primarily based in physical causation, it can be extended to non-physical causation by analogy. However, unlike other physicalist models, the dynamics model does not equate causation with the transfer or exchange of a physical quantity. Rather, it associates causation with a pattern of forces and a position vector that indicates an endstate. Previous researchers have suggested that causation is closely linked to the notion of force (Ahn & Kalish, 2000, Bigelow et al., 1988; Leslie, 1994). In particular, Bigelow & Pargetter (1990) proposed that causation might be associated with a specific pattern of several forces, though they did not specify the exact pattern. Important parts of the dynamics model are also reflected in diSessa's (1993) phenomenological primitives, as well as in White's (2000) influence and resistance model, in which causal judgments are likened to the passage of energy in a physical system.

The importance of force in the representation of causation is illustrated by the causal (but static) situations described in (2).

(2) a. Pressure will cause the water to remain liquid at slightly below 0°C.
b. Dirt caused the valve to stay open.
c. Tiny barbs on the stinger cause it to remain in the wound.
d. Guide wires prevented the Christmas tree from falling.

In each of the situations described in (2) nothing happens. Because nothing happens, there is no regular sequence of events, or transfer or exchange of energy, at least at the macro-level. What is present in each of these situations is a configuration of forces. According to the dynamics model, it is this configuration of forces that makes these situations causal (2a-c) or preventative (2d).

The dynamics model is based on Talmy's (1985, 1988) *force dynamics* account of causation (see also Jackendoff, 1990; Pinker, 1989; Siskind, 2000; Verhagen & Kemmer, 1997; Verhagen, 2002; Wolff, 2003; Wolff, 2007; Wolff & Zettergren, 2002). By analyzing the concept of CAUSE into patterns of forces, Talmy showed that the concept of CAUSE could not only be grounded in properties of the world but also be used to define other related concepts such as

ENABLE, PREVENT, and DESPITE. He also showed how this approach to causation could be extended to many domains of experience, including the physical, intra-psychological, social, and institutional. I incorporate many of Talmy's key ideas into the dynamics model of causation. However, I also introduce several new distinctions and makes significant changes to the theory's underlying semantics.[1]

The dynamics model holds that the concept of CAUSE and related concepts involve interactions between two main entities: an affector and a patient (the entity acted on by the affector). The nature of this interaction can be described at two levels of analysis. The *category level* specifies summary properties of various cause-related concepts. Distinctions at this level are sufficient to distinguish different classes of causal verbs (see Wolff, Klettke, Ventura, & Song, 2005). The *computational level* re-describes the distinctions at the category in terms of units of cognition that represent physical quantities in the world. It is at this level that causes and related concepts are explicitly linked to configurations of force.

The category level of representation. The dynamics model holds that, at the category level, the concept of CAUSE and related concepts can be understood in terms of three dimensions (Wolff & Song, 2003). Specifically, as summarized in Table 1, the concepts of CAUSE, ENABLE, PREVENT, and DESPITE can be captured in terms of 1) the *tendency* of the patient for the endstate, 2) the presence or absence of *concordance* between the affector and the patient, and 3) *progress toward the endstate*.

*********************** Insert Table 1 about here *************************

Table 1. Representations of CAUSE, ENABLE & PREVENT

	Patient tendency for endstate	Affector-patient concordance	Result: endstate approached
CAUSE	N	N	Y
ENABLE	Y	Y	Y
PREVENT	Y	N	N

Note. Y = Yes, N = No

The semantics of these three dimensions are illustrated by the sentences in (3). Consider the example of causation in (3a). In this sentence, the patient (the boat) does not have a tendency for

[1] See Wolff (2006) for summary of the key differences between the two accounts.

the endstate (heeling). The affector (wind) is not in concordance with the patient and the result occurs. In enabling situations, as in (3b), the tendency of the patient (the body) is for the result (to digest food). The affector (vitamin B) does not oppose the patient, and the result occurs. In preventing situations, as in (3c), the patient (the tar) has a tendency for the result (bonding). The affector (the rain) opposes the tendency of the patient and the result does not occur.

(3) a. Wind caused the boat to heel.
b. Vitamin B enables the body to digest food.
c. Rain prevented the tar from bonding.

The computational level of analysis. The computational level of the dynamics model re-describes the three dimensions of tendency, concordance, and result in terms of patterns of forces, or vectors. In discussing such vectors I make a distinction between vectors in the world and vectors in people's minds. Vectors in the world are quantities that have a point of origin, a direction, and a magnitude. The vectors in people's representations of causation are more qualitative. Specifically, vectors in people's representations are predicted to be relatively accurate with respect to direction, but somewhat imprecise with respect to magnitude. People may be able to infer the relative magnitude of two vectors, that one is greater than another. Uncertainty about the magnitude of the vectors adds a certain amount of indeterminacy to people's representations of force dynamic concepts. It is hypothesized that our mental notion of force vectors can represent not only physical forces but also social and psychological forces. Like physical forces, social and psychological forces can be understood as quantities that influence behavior in a certain direction. In this paper, all vectors are in boldface (e.g., **P**).

The dynamics model specifies that four types of force vectors are relevant to the mental representation of cause-related concepts. **A** represents the force exerted on the patient by the affector; **P** represents the force (i.e., thrust) produced by the patient itself or, in the absence of such a force, its weight (e.g., gravity) and/or resistance to motion due to frictional forces; and **O** represents the summation of the remaining other forces acting on the patient. The patient's path through space is determined by **R**, the resultant force acting on the patient based on the vector addition of **A**, **P**, and **O**. In addition to these four forces, people's mental representation of the patient's location with respect to an endstate is specified by the vector **E,** which begins at the patient and ends at the endstate, as shown in Figure 1.

************************ Insert Figure 1 about here **************************

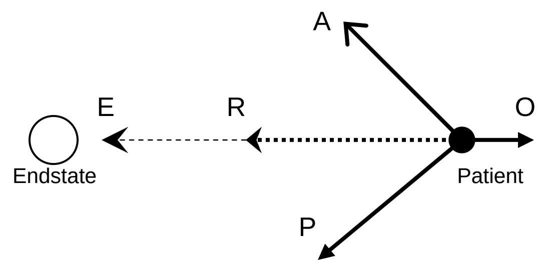

Figure 1. Forces associated with the affector, **A**, patient **P**, and other forces, **O**, combine to produce a resultant force, **R**, that is directed toward the endstate, as specified by the position vector, **E**.

With these definitions and assumptions in place, the relationship between the category and computational levels of the dynamics model can be specified, as summarized in Table 2.

************************ Insert Table 2 about here *************************

Table 2. Dimensions in dynamics model

Tendency (of patient for endstate)	**P** & **E** are collinear
Concordance (of affector & patient)	**A** & **P** are collinear
Result: Endstate approached	**R** & **E** are collinear

Tendency - As shown in Table 2, the patient can be viewed as having a tendency for the endstate when the force associated with it, **P,** is in the direction of the endstate, **E**, that is, when **P** and **E** are collinear. For example, in the free-body diagrams illustrating ENABLE, PREVENT and DESPITE in Figure 2, **P** lies in the same direction as **E**, indicating that the patient has a tendency for the endstate. In the CAUSE configuration, **P** does not point in the same direction as **E**, indicating that the patient does not have a tendency for the endstate.

Concordance - The patient and the affector are in concordance when the patient's tendency, **P**, is in the same direction as the force associated with the affector, **A**, that is, when **P** and **A** are collinear. As shown in Figure 2, collinearity holds in the case of ENABLE but not in the cases of CAUSE, PREVENT, and DESPITE.

Result – The patient will approach the endstate and eventually reach it, barring changes in the forces acting on the patient, when the sum of the forces acting on the patient, **R**, is in the direction of the endstate **E**, that is, when **R** and **E** are collinear.

************************ Insert Figure 2 about here **************************

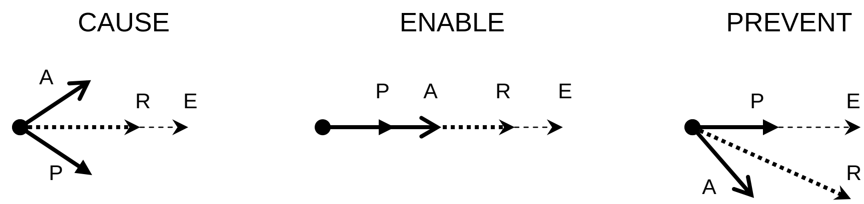

Figure 2. Configurations of forces associated with CAUSE, ENABLE, and PREVENT.

Spanning restriction and heuristic. The dynamics model places constraints on what constitutes a valid configuration. Valid configurations are those in which the resultant could be produced from the vector addition of the component vectors. Thus, according to the dynamics model, understanding causal relationships involves evaluating whether **R** reflects the sum of the vectors **A**, **P**, and **O** in the real world. The model assumes that people are sensitive to the way in which forces interact in the real world. However, since vectors in the mind do not have exact magnitudes, their representations do not allow for exact vector addition to assess **R**. Instead of exact vector addition, I propose that people use a qualitative criterion for deciding whether a resultant could have been produced from the vector addition of two vectors. An implication of the parallelogram law of vector addition is that the resultant of two vectors will always lie on top of or within the region, or *span*,[2] bounded by the vectors being added, as depicted in Figure 3.

[2] The word "span" is used here in a more restricted sense than is used in mathematics. In its usual sense, "span" refers to, for example, the set of resultant vectors, $\mathbf{u}_i$, that can be formed from the equation $\mathbf{u} = c_1\mathbf{v}_1 + c_2\mathbf{v}_2$, where $\mathbf{v}_1$ and $\mathbf{v}_2$ are vectors and c_1 and c_2 are scalars. When using "span" in the context of the dynamics model, I restrict c_1 and c_2 to values that are equal to or greater than zero, thus limiting the resultant vectors, $\mathbf{u}_i$, to the region bounded by and including $\mathbf{v}_1$ and $\mathbf{v}_2$.

************************ Insert Figure 3 about here *************************

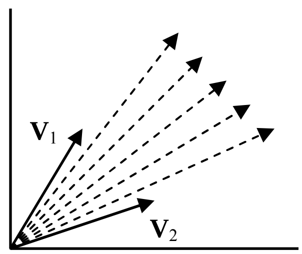

Figure 3. Despite uncertainty about the magnitudes of $\mathbf{V}_1$ and $\mathbf{V}_2$, we can infer that the resultant of the two vectors will reside within the area bounded by $\mathbf{V}_1$ and $\mathbf{V}_2$.

If the resultant lies outside the span of the two vectors being added, the configuration violates the *spanning restriction*. According to the dynamics model, people refer to the spanning restriction in a heuristic—the *spanning heuristic*—to make rough guesses about whether a resultant was produced from the vector addition of the component forces. When a resultant—as indicated by a patient's motion— lies within the span bounded by two vectors, the spanning heuristic warrants the inference that the resultant was produced from the vector addition of the two component vectors. When a resultant lies outside the span, the spanning heuristic holds that the result was not due to the addition of the two component vectors alone.

Testing the dynamics model

The dynamics model makes predictions about how people will interpret scenes that instantiate different configurations of forces. In particular, the dynamics model makes predictions about what kinds of events will count as causation, as opposed to enablement or prevention. Here I provide a summary of the experiments reported in Wolff & Zettergren (2002) and Wolff (2007). In these experiments, participants viewed 3D animations that depicted a boat (the patient) moving in a pool of water toward a cone (the endstate); a bank of fans beside the pool (the affector) affected the speed and direction of the boat (see Figure 3). The boat's path through the water was completely determined by a physics simulator, which allowed us to independently manipulate the dynamics of the situation, in particular, the forces acting on the boat. After viewing an animation, participants chose from several linguistic descriptions or "none of the above" (Wolff, 2007). All of the descriptions were the same (*The fans ____ the boat to [from] hit[ting] the cone*) except for the main verb, which was either *caused*, *helped* or *prevented*.

************************ Insert Figure 4 about here **************************

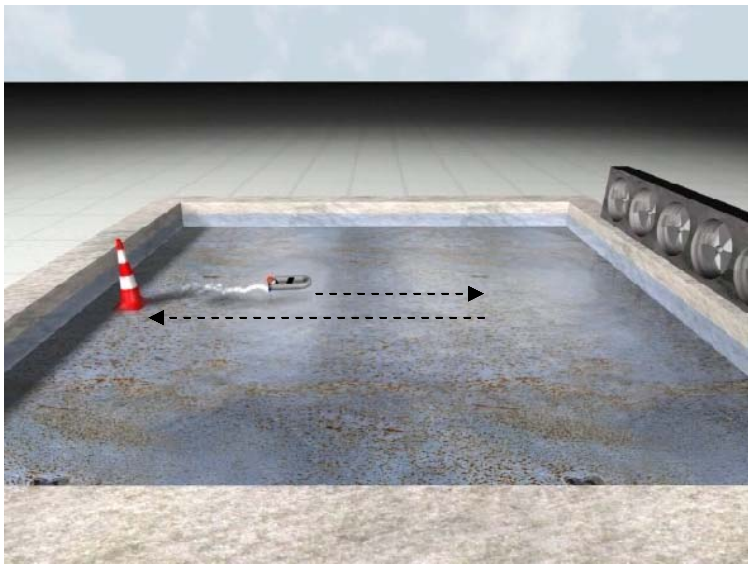

Figure 4: Frame from an animation used to instantiate a CAUSE interaction.

We predicted that CAUSE descriptions would be chosen when the boat initially moved away from the cone (Tendency = N), but eventually hit it because of the fans' blowing in the direction of the cone (Concordance = N; Endstate approached = Y). We predicted that ENABLE descriptions would be chosen when the boat moved toward the cone (Tendency = Y), the fans blew in the same direction (Concordance = Y), and the boat ultimately reached the cone (Endstate approached = Y). We predicted that PREVENT descriptions would be chosen when the boat moved toward the cone (Tendency = Y), but the fans blew in another direction (Concordance = N) such that the boat missed the cone (Endstate approached = N). Finally, we predicted participants would choose the option "none of the above" when none of the above configurations were instantiated.

The predictions of the dynamics model were fully borne out by the results. The lower portion of Table 3 shows the percentage of times people chose each of the four options for each of the vector configurations. Importantly, participants did not choose *prevent* for all scenes in which the boat missed the cone (6-8). Instead, *prevent* was restricted to just those situations in which the boat had an initial tendency for the endstate (4). Likewise, participants did not choose *cause* or *help* just because the boat hit the cone (5). Since two vectors must be evaluated to determine concordance, this last result strongly suggests that participants considered the relationships among the vectors when choosing the description, just as predicted by the theory.

*********************** Insert Table 3 about here ************************

Table 3. Experiment 1 in Wolff (2007), along with associated predictions and results

Config. #	1	2	3	4	5	6	7	8
Affector ( ) Patient ( ) Result. ( )	E	E	E	E	E	E	E	E
Predicted	CAUSE	ENABLE	ENABLE	PREVENT	Unspecified	Unspecified	Unspecified	Unspecified
"Cause"	**94%**	11%	6%	-	-	-	6%	-
"Help"	6%	**89%**	**94%**	-	11%	-	-	-
"Prevent"	-	-	-	**100%**	6%	-	-	6%
"No verb"	-	-	-	-	**83%**	**100%**	**94%**	**94%**

The dynamics model easily extends beyond the one-dimensional interactions used in Experiment 1 to two-dimensional interactions. Wolff (2007; see also Wolff & Zettergren, 2002) demonstrated this using the configurations shown in Table 4. Also as shown in Table 4, participants' choices matched the predictions of the theory.

*********************** Insert Table 4 about here ************************

Table 4. Configurations used in Experiment 2 of Wolff (2007), with associated predictions and results

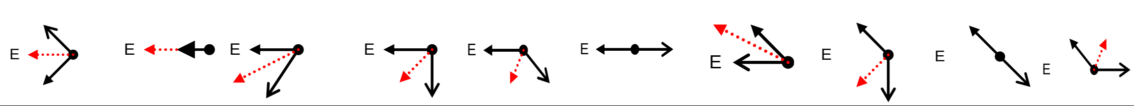

Config. #	1	2	3	4	5	6	7	8	9	10
Affector ( ) Patient ( ) Result. ( )	E	E	E	E	E	E	E	E	E	E
Predicted	CAUSE	ENABLE	PREVNT	PREVNT	PREVNT	PREVNT	Unspecified	Unspecified	Unspecified	Unspecified
"Cause"	**89%**	11%	-	-	-	-	-	-	-	-
"Help"	.1%	**83%**	-	-	-	-	-	-	-	-
"Prevent"	-	-	**94%**	**94%**	**89%**	**89%**	-	17%	-	11%
"No verb"	-	6%	6%	6%	11%	11%	**100%**	**83%**	**100%**	**89%**

The results from Wolff (2007) are consistent with the hypothesis that people's causal concepts are based on configurations of force. The results also support the dynamics model's account of how people determine causation on the basis of a single observation. According to the model, people identify causal relationships by constructing representations of the forces acting on the patient. However, the data so far are open to an alternative possibility; specifically, they could be explained in terms of kinematics rather than dynamics. In a kinematics account, only visible movements—specifically, the velocities—are considered in the classification of interactions. For example, causation might be defined as an interaction in which the patient was not moving toward the endstate, but then moved toward the endstate once the affector made contact with it. Enablement could be defined as an interaction in which the patient was moving toward the endstate, but then moved more quickly toward the endstate once the affector made contact with it. Finally, prevention might be defined as an interaction in which the patient was moving toward the endstate, but then moved away from the endstate once the affector made contact with it.

One way to test between kinematics and dynamic approaches to causation would be to examine whether people are aware of the way in which forces are added. If people's causal judgments are based on kinematics, peoples' causal judgments should be insensitive to such violations in the way the forces are added. On the other hand, if causal judgments are based on the dynamics of an event, people should notice when an object moves in a way that is not consistent with the way forces are added.

As discussed earlier, it is assumed that people use a qualitative criterion, *the spanning heuristic*, to determine whether a particular resultant could be derived from a particular set of forces. When a patient moves in a direction that lies within the area between the forces acting on the patient (see Figure 4), the spanning heuristic should lead people to assume that the resultant is produced from the vector addition of those forces. Conversely, when the resultant does not reside within the span of the component vectors, it can be said that the configuration violates the *spanning restriction*.

The spanning heuristic provides a rough method of evaluating whether the net force acting on a patient is derivable from the overt forces acting on the patient. However, in certain circumstances, the heuristic may lead people to incorrectly infer that the net force acting on the patient is fully explained in terms of the perceived forces when, in fact, there are other forces in play. Such an illusion of sufficiency is most likely to occur when there is more than one external force acting on the patient, that is, when the magnitude of other forces, **O,** is greater than 0. For example, consider the three scenes and free-body diagrams in Figure 5. The forces entered into the physics simulator for all the scenes are depicted in the first free-body diagram. In the first animation, a boat motors to the middle of a pool, two sets of fans turn on, and the boat moves toward the cone and ultimately hits it. The second panel shows a frame from an animation that is exactly the same as the one on the left except that the one of the fans is not shown (though its force is still present). In this animation, the boat moves into the area bounded by the overt forces; hence, according to the spanning heuristic, the fan may be construed as a cause of the boat's hitting the cone. The third panel shows an animation that is also exactly the same as the one in the first panel except that the opposite fan is not shown. In this scene, based on single visible fan, the boat's direction lies outside the area bounded by the perceived forces. According to the spanning heuristic, then, the visible fan cannot be construed as a cause of the boat's hitting the cone.

************************ Insert Figure 5 about here *************************

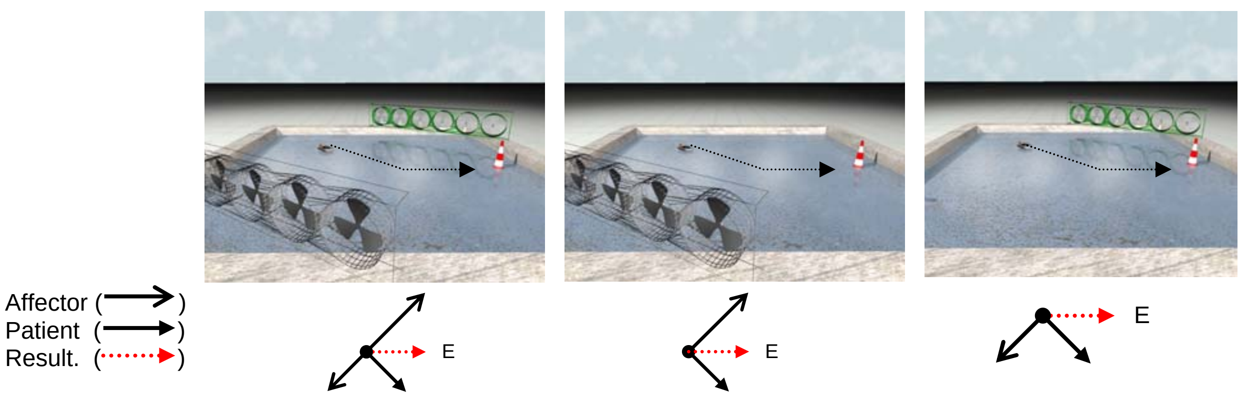

Figure 5. In each animation, the boat motors to the middle, the fans turn on, the boat changes course, and the boat hits the cone. Each animation is based on the same configuration of forces as shown in the first panel. However, in the second and third panels, only one of the two fans appears in the animation, as implied by the incomplete.

These predictions were tested in an experiment in which participants saw four pairs of animations (Wolff, 2006). Like the middle and rightmost animations in Figure 5, these animations depicted situations in which two external forces were in play, but only one was shown. One member of each minimal pair depicted a situation in which the resultant was within the span of the observable forces and the other member depicted a situation in which the resultant was not within that span. As predicted, participants were quite willing to say that the fan “caused” the boat to hit the cone when the resulting direction of the boat was within the span of the two observable forces (percent “caused” = 84%). Also, as predicted, participants were quite unwilling to say that the fan “caused” the boat to hit the cone when the boat moved in a direction that was outside of the area spanned by the two observable forces (percent caused = 18%). When the boat moved outside of the area span of the observable forces, people decided that it was not physically possible for the result to be due to the observable forces alone.

In both kinds of scenarios, the boat changed direction immediately after the fans began to blow, so neither temporal nor spatial contiguity was a determining factor in participants’ judgments. The two kinds of animations were equally “physically accurate” since they were based on exactly the same underlying forces. They were also equally “incomplete” in that only one of the two forces acting on the boat was shown. Finally, both kinds of animations were equally “natural” in the sense that we are likely to encounter scenes in which any of a number of external forces may be hidden.

Support for the spanning heuristic also makes clear that that people’s judgments of causation do not require knowing the exact magnitudes of the forces. In the spanning conditions, the animations did not provide enough information to determine whether the boat’s course was due to the force associated with just the one fan or due to that force in combination with (an)other hidden force(s). Nevertheless, when the boat moved within the span of the overt forces, people agreed that the fan “caused” the boat to hit the cone. Thus, precise knowledge of magnitudes is not necessary for classifying the situation as causal. What appears to be necessary, instead, is awareness of the direction of the forces, which supports the hypothesis that people think about causation in terms of vectors. In sum, the results of this experiment support three main assumptions of the dynamics model, namely, that people think about causal situations in terms of vectors, that they perform vector addition via the spanning heuristic, and that their causal judgments are based on the dynamic properties of an event, which are derived from the kinematics.

In addition to these results, several other problems remain for a kinematics-only account of causation. As discussed above, the concept of CAUSE extends to situations in which there are conflicting forces, but no change occurs (e.g., *The rubber bottom caused the cup to stay in place*). A kinematics-based account cannot distinguish these static causal situations from situations that are static but simply spatial (e.g., **The tree causes the roof to be under the branch*). Another limitation to a kinematics approach is that it does not easily explain our language for non-physical causation. In describing social causation, we rarely talk about "social velocities" or "peer accelerations." Rather, we talk about "social forces" and "peer pressures." Our language for describing social interactions implies that we think about these interactions in terms of dynamic properties. We can talk about "psychological forces" as when we describe someone as strong willed or internally conflicted. Ordering someone to do something is easily viewed as imparting an invisible, directional influence that may or may not produce an effect.

These patterns in language suggest that the dynamics model might be extended to account for people's causal judgments of scenes involving intentions and desires. Consider, for example, the scenarios depicted in Figure 6, in which a woman is standing in a raft and pointing in a particular direction. She indicates the direction she wants to move by pointing. If intentions are analogous to physical forces, as assumed in the dynamics model, people should prefer to say that the fans *caused* the woman to reach the cone when the woman is pointing away from the cone and that the fans *enabled* her to reach the cone when the woman is pointing towards the cone. In addition, they should report that the fans *prevented* her from reaching the cone when she points towards the cone, but is pushed away from it.

************************ Insert Figure 6 about here *************************

CAUSE

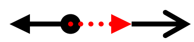

Figure 6. The scene depicts a CAUSE situation since the woman does not want to go to the cone (as indicated by the direction of her pointing) but the fans push her there nevertheless.

These possibilities were tested in an experiment involving intentional forces (Wolff, 2006). Participants saw two kinds of animations. Half of the animations were based on physical forces only, just as in the previous experiments. In the remaining animations, the underlying configurations of forces were exactly the same as those used in the physical-only animations. The main difference was that the patient’s tendency was indicated by a woman’s intention (as represented by her pointing). Specifically, in physical-and-intentional animations, the motor boat was replaced with a round rubber raft with a woman inside of it. Because the boat did not have an engine, the tendency of the patient (i.e., the woman in the round rubber raft) could not be specified by its self-motion. Rather, for these animations, the patient’s tendency was specified by the woman pointing in a certain direction. Thus, for the CAUSE animations, the tendency of the patient was indicated by the boat motoring away from the cone in the physical force-only condition and by the woman pointing away from the cone in the physical-and-intentional force condition. In the ENABLE and PREVENT animations, the tendency of the patient was indicated by the boat motoring toward the cone and the woman pointing toward the cone.

As in the previous experiments, participants saw the animations and described them by choosing sentences containing either cause-verbs (*cause, get, make*), enable-verbs (*enable, help, let*), prevent-verbs (*block, keep, prevent*), or the option "none of the above." The results showed that people treated the woman's intention as if it were a physical force. In the animations in which both forces were physical, the results were the same as in previous experiments: people prefer cause-verbs for the CAUSE animations (94%), enable-verbs for the ENABLE animations (89%), prevent-verbs for the PREVENT animations (100%). The pattern of responses was the same when the patient's tendency was determined by where the woman was pointing. People preferred cause-verbs for the CAUSE animations (94%), enable-verbs for the ENABLE animations (83%), prevent-verbs for the PREVENT animations (100%). The results indicated that the dynamics model can be extended to situations involving non-physical forces. In Wolff (2006), I show that the dynamics model extends to situations in which all of the forces are either intentions or desires.

Summary

The dynamics model describes how people's representations of causation reproduce causation in the real world. It explains how people's representations of causation are specified in terms of forces, which are the very quantities that cause events in the real world. The assumptions embodied in the dynamics model contrast with those of Hume and models of the Humean tradition. Hume (1748/1975) maintained that people's conceptualization of causation went beyond kinematics to include such notions as force, necessary connection, causal power and/or energy. Importantly, however, he also held that because such notions could not be directly seen, they were constructs of the mind rather than reflections of what went on in the world. Because these quantities could not be seen, Hume argued, they must emerge from the observation of a regular succession of events.

Michotte's (1946/1963) assumptions about causation in launching events were very similar to Hume's assumptions about causation in general. Like Hume, Michotte argued that people's notion of causation extended beyond the kinematics of an event: when an object A strikes an object B and sends it into motion, people inferred the presence of a causal relationships, not just a sequence of motions. Also like Hume, Michotte felt that the impression of causation did not directly reflect the dynamics of the event. Instead, he felt that the impression was a visual illusion formed from a (possibly innate) perceptual mechanism. Michotte's view of causation differed from Hume's in that he held that the causal impression could be formed on the basis of a single observation rather than requiring multiple occurrences for its apprehension.

The dynamics model shares certain assumptions with Michotte's account of causation for launching events. Both accounts hold that people's representations of causation extend beyond kinematics and that the induction of causation can occur on the basis of a single observation. However, unlike Michotte's and Hume's accounts, the dynamics model holds that people's representations of causation copy, or replicate, certain aspects of the dynamics of an event, and as a consequence, people's causal representations capture important aspects of the quantities that bring about causation in the world. This does not imply that people are able to recover all of the dynamic properties of an event. Work in physics understanding shows that this clearly does not happen. However, people do appear to be able to construct partial representations of an event's dynamics. In particular, they appear to be relatively good at recovering the direction of forces.

Launching events and dynamics

Given the close similarities between the dynamics model and Michotte's account of causation, one might wonder why Michotte chose to argue against a dynamics account of the causal impression. As reviewed below, the evidence that Michotte used to argue against a dynamics account of causation is open to alternative explanations. In addition, when we look more closely at Michotte's own experiments, we find that many of his results provide evidence against his own hypotheses and support for dynamics.

Michotte's arguments against dynamics. One reason why Michotte felt that the causal impression was not due to dynamics is that people sometimes reported perceiving causation in events that he viewed as physically impossible. For example, people reported perceiving causation in situations in which objects A and B are moving, object A faster than object B, and when the two make contact, object A stops and object B slows down (Michotte, 1946/1963, p. 71). On the basis of such results, Michotte concluded that the causal impression is not based on past experience with the world. Had people referred to past experience, they would have expected object B to speed up, not slow down, after being hit. However, while such a sequence of events may be unusual, it is not necessarily at odds with Newtonian physics. Friction can change dramatically over the course of an object's movement, as when a ball rolls off an asphalt road and onto a gravel driveway. Michotte's "impossible" event is not, in fact, impossible in the world, and so his finding does not necessarily rule out the role of dynamics in the perception of causation.

Another of Michotte's arguments for the independence of the launching effect and real world causation is that sometimes the causal impression failed to obtain for trajectories that people experience in the real world. In support of this point, Michotte conducted several experiments (34 and 35) in which object A hits object B directly, and B travels at an angle away from its expected straight line path. The degree of deviation from B's expected straight path ranged from 25° to 90°; as the size of the angle increased, the impression of causation grew weaker. Michotte points out that this result is at variance with our real world experience in which two colliding objects can travel at angles (besides 180°) and still be viewed as causal (e.g., billiards, marbles). However, Michotte's collision events were quite different from those involving billiard balls and marbles. In particular, since object A hit object B directly, Newtonian physics would predict that object B should move straight ahead, not at an angle. In addition, in Michotte's Experiment 35, from which it seems most of his conclusions were derived, the objects were rectangles instead of circles. If a rectangle were to hit another rectangle straight on, it would, indeed, be quite surprising and inconsistent with Newtonian physics if the second rectangle veered off at an angle, especially a right angle. Thus, in complete contrast to Michotte's stated conclusions, his work showing how changes in direction weaken the causal impression actually supports the conclusion that there is a tight relationship between the causal impression and dynamics in the world.

Yet another reason why Michotte believed the causal impression was only a perceptual phenomenon is that people experience the causal impression even when the objects involved are spots of light, shadows, or lines painted on a rotated disk. In other words, people perceive causation while also knowing that such causation does not occur in the real world (1946/1963, p. 84-85). However, a dynamics approach to causation does not imply that people cannot be subject to illusions of causation. A particular configuration of forces will produce only one kinematic pattern, but a single kinematic pattern is potentially consistent with more than one configuration of forces. This asymmetry explains why causal illusions can sometimes occur: people may infer the wrong configuration of forces from a particular kinematic pattern. This is especially likely when the actual forces driving the kinematics are obscured, as in the case of Michotte's launching events. Further, the process of inducing forces is likely to be at least partially automatic (Runeson & Frykolm, 1983), so causal illusions may occur even when the inferred configuration of forces is inconsistent with prior knowledge of the situation.

How Michotte's findings indicate the role of dynamics in the perception of causation. Michotte emphasized that the causal impression was not a mere copy or reproduction of what goes on in the real world. If anything, however, many of Michotte's findings indicate just the opposite. For example, Michotte observed that the causal impression disappeared fully when there was a gap of around 150 ms between the moment objects A and B made contact and the moment B began to move. This finding is readily explained by the dynamics model. When object A hits object B, the force imparted on B is instantaneous. If object B begins moving well after it is hit, its movement cannot be due to the force imparted by object A. Another finding of Michotte's is that the perception of causation is strongest when object A makes physical contact with object B. This finding is also consistent with a dynamics approach since contact forces cannot exist unless objects make contact with one another.

Yet another finding of Michotte's concerns a phenomenon he referred to as the *radii of action*. The radii of action are the portions of the paths traveled by A and B that subjectively appear to be relevant to the impression of causation. In particular, when B travels beyond A's radius of action, it appears to be moving on its own, not as a consequence of the collision. Michotte found that object B's radius of action increased with the speed of object A. Michotte offered no explanation for the phenomenon of radii of action. However, the dynamics model offers a natural explanation for this effect: as object A's speed increases, the force it imparts on B increases, and in turn, so will the distance B travels as a consequence of the impact of A (for a related proposal, see Hubbard & Ruppel, 2002).

Finally, according to Michotte, the causal impression should be strongest when the two parts of a launching event constitute a single continuous movement, whereby the motion of the first object extends into the second and there is an "ampliation of motion." According to this hypothesis, any differences in velocity between the first and second objects should decrease the causal impression since any differences in velocity would make the sequence of events less continuous. However, in contrast to this prediction, Michotte found that the causal impression was stronger when the speed of object B was slower than that of object A. Specifically, in Experiments 15 and 39, people reported a much stronger causal impression when the ratio in speed of objects A and B was 4:1 than when the ratio was 1:1. This result is consistent with the dynamics model, which predicts that the second object should move less rapidly than the first because the second object resists moving in the direction of the force acting on it. That resistance, due to friction and inertia, means that the second object will move more slowly than the first (and ultimately will come to a stop). When object B's speed is the same as object A's, the dynamics model predicts that the causal impression should be weaker because of the absence of evidence for such resistance.

Conclusion

Humean theories of causation do not deny that people's everyday notions of causation are associated with invisible quantities such as energy or force. But they would say that such notions were better viewed as part of the occult than the actual world. For these theories, *out of sight* is *out of mind*.

A physicalist approach does not deny the importance of kinematic features, but such features do not form the basis for people's causal representations; rather, they are the keys for unlocking the dynamics of an event. People's ability to infer these properties is by no means perfect, but neither is it arbitrary. In particular, people may be reasonably good at inferring the presence of forces and their direction, but relatively insensitive to their magnitude. The dynamics of an event are central to people's concept of causation because they are central to causation in the actual world. Because dynamic properties can be sensed, a physicalist approach to causation not only grounds causation in the world, it also explains how causation might be experienced in our own bodies, and why such notions of causal power, energy, and force are not just side-effects of statistical dependencies.

References

Ahn, W., & Bailenson, J. (1996). Mechanism-based explanations of causal attribution: An explanation of conjunction and discounting effect. *Cognitive Psychology, 31*, 82-123.

Ahn, W., & Kalish, C. W. (2000). The role of mechanism beliefs in causal reasoning. In F. C. Keil & R. A. Wilson (Eds.), *Explanation and cognition* (pp. 199-225). Cambridge, MA: The MIT Press.

Ahn, W., Kalish, C. W., Medin, D. L., & Gelman, S. A. (1995). The role of covariation versus mechanism information in causal attribution. *Cognition, 54*, 299-352.

Aronson, J. L. (1971). On the grammar of 'CAUSE'. *Synthese, 22*, 414-430.

Bigelow, J., & Pargetter, R. (1990). Metaphysics of causation. *Erkenntnis, 33*, 89-119.

Bigelow, J., Ellis, B., & Pargetter, R. (1988). Forces. *Philosophy of Science, 55*, 614-630.

Bullock, M., Gelman, R., & Baillargeon, R. (1982). The development of causal reasoning. In W. Friedman (Ed.), *The developmental psychology of time* (pp. 209-255). London: Academic Press.

Bunge, M. (1959). *Causality: the place of the causal principle in modern science*. Cambridge, MA, Harvard University Press.

Cheng, P. W. (1997). From covariation to causation: A causal power theory. *Psychological Review, 104*, 367-405.

Cheng, P. W., & Novick, L. R. (1991). Causes versus enabling conditions. *Cognition, 40*, 83-120.

Cheng, P. W., & Novick, L. R. (1992). Covariation in natural causal induction. *Psychological Review, 99*, 365-382.

Clement, J. (1983). A conceptual model discussed by Galileo and used intuitively by physics students. In D. Gentner & A. L. Stevens (Eds.), *Mental models* (pp. 299-324). Hillsdale, NJ: Lawrence Erlbaum Assoc.

Cohen, L. B., Amsel, G., Redford, M. A., & Casasola, M. (1998). The development of infant causal perception. In A. Slater (ed.), *Perceptual development: Visual, auditory, and speech perception in infancy* (pp. 167-209). East Sussex, UK: Psychology Press Ltd.

Choi, H. & Scholl, B. J. (2004). Effects of grouping and attention on the perception of causality. *Perception & Psychophysics, 66*, 926-942.

diSessa, A. (1993). Towards an epistemology of physics. *Cognition and Instruction, 10*, 105-225.

Dowe, P. (2000). *Physical causation*. Cambridge University Press: Cambridge, UK.

Fair, D. (1979). "Causation and the flow of energy." *Erkenntnis 14*, 219-250.

Gilden, D. L. (1991). On the origins of dynamical awareness. *Psychological Review, 98*, 554-568.

Goldvarg, E., & Johnson-Laird, P. (2001). Naive causality: A mental model theory of causal meaning and reasoning. *Cognitive Science, 25*, 565-610.

Gopnik, A., Glymour, C., Sobel, D., Shulz, L., Kushnir, T., & Danks, D. (2004). A theory of causal learning in children: Causal maps and Bayes nets. *Psychological Review, 111*, 1-31.

Hecht, H. (1996). Heuristics and invariants in dynamic event perception: Immunized concepts or nonstatements? *Psychonomic Bulletin & Review, 3*, 61-70.

Hume, D. (1748/1975). *Enquiry concerning Human Understanding*, in *Enquiries concerning Human Understanding and concerning the Principles of Morals*, edited by L. A. Selby-Bigge, 3rd edition revised by P. H. Nidditch, Oxford: Clarendon Press.

Hume, D. (1737/1978). *A Treatise of Human Nature*, edited by L. A. Selby-Bigge, 2nd edition revised by P. H. Nidditch, Oxford: Oxford University Press.

Jackendoff, R. (1990). *Semantic structures*. Cambridge, MA: The MIT Press.

Joskowicz, L., & Sacks, E. (1991). Computational kinematics. *Artificial Intelligence, 51*, 381-416.

Kaiser, M. K., Proffitt, D. R., Whelan, S. M., & Hecht, H. (1992). The influence of animation on dynamical judgments. *Journal of Experimental Psychology: Human Perception and Performance, 18*, 669-690.

Leslie, A.M. (1994). ToMM, ToBy, and agency: core architecture and domain specificity. In L. Hirschfield, & S. Gelman (eds), *Mapping the Mind: Domain Specificity in Cognition and Culture* (119–148). Cambridge University Press.

Luhmann, C. C., & Ahn, W. (2005). The meaning and computation of causal power: Comment on Cheng (1997) and Novick and Cheng (2004). *Psychological Review, 112*, 685-692.

Mackie, J. L. (1974). *The cement of the universe*. Oxford: Oxford University Press.

McCloskey, M. (1983). Naïve theories of motion. In D. Gentner & A. L. Stevens (Eds.), *Mental models* (pp. 299-324). Hillsdale, NJ: Lawrence Erlbaum Associates.

McCloskey, M., & Kohl, D. (1983). Naïve physics: The curvilinear impetus principle and its role in interactions with moving objects. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 9*, 146-156.

McCloskey, M., Washburn, A., & Felch, L. (1983). Intuitive physics: The straight-down belief and its origin. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 9*, 636-649.

Michotte, A. E. (1946/1963). *The perception of causality*. New York: Basic Books.

Pearl, J. (2000). *Causality: Models, reasoning, and inference*. Cambridge, UK: Cambridge University Press.

Pinker, S. (1989). *Learnability and cognition: The acquisition of argument structure*. Cambridge, MA: The MIT Press.

Pittenger, J. B. (1990). Detection of violations of the law of pendulum motion: Observers' sensitivity to the relation between period and length. *Ecological Psychology, 2*, 55-81.

Proffitt, D. R., & Gilden, D. L. (1989). Understanding natural dynamics. *Journal of Experimental Psychology: Human Perception and Performance, 15*, 384-393.

Runeson, S., & Frykholm, G. (1983). Kinematic specification of dynamics as an informational basis for person and action perception: Expectation, gender recognition, and deceptive intention. *Journal of Experimental Psychology: General, 112*, 585-615.

Runeson, S., Juslin, P., & Olsson, H. (2000). Visual perception of dynamic properties: cue heuristic versus direct-perceptual competence. *Psychological Review, 107*, 525-555.

Runeson, S., & Vedeler, D. (1993). The indispensability of precollision kinematics in the visual perception of relative mass. *Perception & Psychophysics, 53*, 617-632.

Russell, B. (1948). *Human Knowledge.* New York: Simon and Schuster.

Salmon, W. (1994). Causality without counterfactuals. *Philosophy of Science 61*, 297-312.

Salmon, W. (1998). *Causality and explanation*. Oxford: Oxford University Press.

Scholl, B. J., & Tremoulet, P. D. (2000). Perceptual causality and animacy. *Trends in Cognitive Sciences, 4*, 299-309.

Schwartz, D. L. (1999). Physical imagery: Kinematic versus dynamic models. *Cognitive Psychology, 38*, 433-464.

Shultz, T. R. (1982). Rules of causal attribution. *Monographs of the Society for Research in Child Development, 47*, 1-51.

Siskind, J. M. (2000). Visual event classification via force dynamics. In *Proceedings of the American Association for Artificial Intelligence* (pp. 149-155).

Sloman, S. (2005). *Causal models: How people think about the world and its alternatives*. Oxford: Oxford University Press.

Sloman, S. A., & Lagnado, D. A. (2005). Do we “do”? *Cognitive Science, 29*, 5-39.

Sobel, D. M., Tenenbaum, J. B., & Gopnik, A. (2004). Children's causal inferences from indirect evidence: Backwards blocking and Bayesian reasoning in preschoolers. Cognitive Science, 303-333.

Talmy, L. (1985). Force dynamics in language and thought. In W. Eilfort, P. Kroeber & K. Peterson (Eds.), *Papers from the parasession on causatives and agentivity at the 21st regional meeting, Chicago Linguistics Society* (pp. 293-337). Chicago: Chicago Linguistics Society.

Talmy, L. (1988). Force dynamics in language and cognition. *Cognitive Science, 12*, 49-100.

Tenenbaum, J. B., & Griffiths, T. L. (2001). Structure learning in human causal induction. In T. Leen, T. Dietterich & V. Tresp (Eds.), *Advances in Neural Information Processing Systems 13* (pp. 59-65). Cambridge, MA: MIT Press.

Thinés, G., Costal, A., & Butterworth, G. (eds.) (1991). *Michotte’s experimental phenomenology of perception*. Hillsdale, NJ: Erlbaum.

Twardy, C. R., & Bingham, G. P. (2002). Causation, causal perception, and conservation laws. *Perception & Psychophysics, 64*, 956-968.

Verhagen, A. (2002). Interpreting usage: Construing the history of Dutch causal verbs. In M. Barlow & S. Kemmer (Eds.), *Usage-based models of language*. Stanford, CA: CSLI Publications.

Verhagen, A., & Kemmer, S. (1997). Interaction and causation: A cognitive approach to causative constructions in Modern Standard Dutch. *Journal of Pragmatics, 27*, 61-82.

White, P. A. (1999). Toward a causal realist account of causal understanding. *American Journal of Psychology, 112*, 605-642.

White, P.A. (2000). Naive analysis of food web dynamics: A study of causal judgment about complex physical systems. *Cognitive Science, 24*, 605-650.

Wolff, P. (2007). Representing causation. *Journal of Experiment Psychology: General*, 136.

Wolff, P. (2003). Direct causation in the linguistic coding and individuation of causal events. *Cognition, 88*, 1-48.

Wolff, P., & Song, G. (2003). Models of causation and the semantics of causal verbs. *Cognitive Psychology, 47*, 276-332.

Wolff, P., & Zettergren, M. (2002). A vector model of causal meaning. In W. D. Gray & C. D. Schunn (Eds.), *Proceedings of the twenty-fourth annual conference of the Cognitive Science Society* (pp. 944-949). Mahwah, NJ: Lawrence Erlbaum Associates.

Table 1. Representations of CAUSE, ENABLE & PREVENT

	Patient tendency for endstate	Affector-patient concordance	Result: endstate approached
CAUSE	N	N	Y
ENABLE	Y	Y	Y
PREVENT	Y	N	N

Note. Y = Yes, N = No

Table 2. Dimensions in dynamics model

Tendency (of patient for endstate)	**P** & **E** are collinear
Concordance (of affector & patient)	**A** & **P** are collinear
Result: Endstate approached	**R** & **E** are collinear

Table 3. Experiment 1 in Wolff (2007), along with associated predictions and results

Config. #	1	2	3	4	5	6	7	8
Affector ( ) Patient ( ) Result. ( )	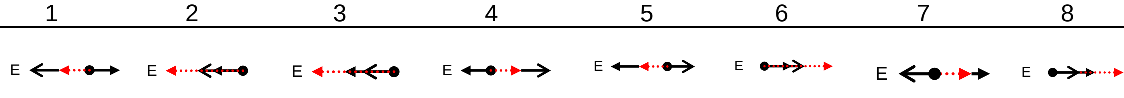							
Predicted	CAUSE	ENABLE	ENABLE	PREVENT	Unspecified	Unspecified	Unspecified	Unspecified
"Cause"	**94%**	11%	6%	-	-	-	6%	-
"Help"	6%	**89%**	**94%**	-	11%	-	-	-
"Prevent"	-	-	-	**100%**	6%	-	-	6%
"No verb"	-	-	-	-	**83%**	**100%**	**94%**	**94%**

Table 4. Configurations used in Experiment 2 of Wolff (2007), with associated predictions and results

Config. #	1	2	3	4	5	6	7	8	9	10
Affector ( ) Patient ( ) Result. ( )	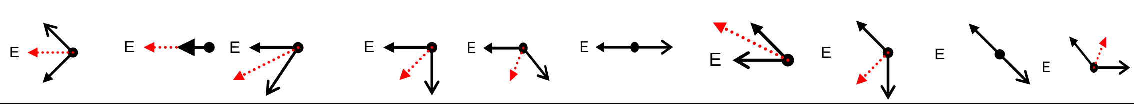									
Predicted	CAUSE	ENABLE	PREVNT	PREVNT	PREVNT	PREVNT	Unspecified	Unspecified	Unspecified	Unspecified
"Cause"	**89%**	11%	-	-	-	-	-	-	-	-
"Help"	.1%	**83%**	-	-	-	-	-	-	-	-
"Prevent"	-	-	**94%**	**94%**	**89%**	**89%**	-	17%	-	11%
"No verb"	-	6%	6%	6%	11%	11%	**100%**	**83%**	**100%**	**89%**

Figure Captions

Figure 1. Forces associated with the affector, **A**, patient **P**, and other forces, **O**, combine to produce a resultant force, **R**, that is directed toward the endstate, as specified by the position vector, **E**.

Figure 2. Configurations of forces associated with CAUSE, ENABLE, PREVENT and DESPITE

Figure 3. Despite uncertainty about the magnitudes of $\mathbf{V}_1$ and $\mathbf{V}_2$, we can infer that the resultant of the two vectors will reside within the area bounded by $\mathbf{V}_1$ and $\mathbf{V}_2$.

Figure 4: Frame from an animation used to instantiate a CAUSE interaction.

Figure 5. In each animation, the boat motors to the middle, the fans turn on, the boat changes course, and the boat hits the cone. Each animation is based on the same configuration of forces as shown in the first panel. However, in the second and third panels, only one of the two fans appears in the animation, as implied by the incomplete.

Figure 6. The scene depicts a CAUSE situation since the woman does not want to go to the cone (as indicated by the direction of her pointing) but the fans push her there nevertheless.

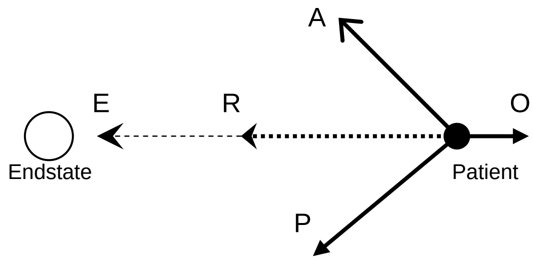

Figure 1. Forces associated with the affector, **A**, patient **P**, and other forces, **O**, combine to produce a resultant force, **R**, that is directed toward the endstate, as specified by the position vector, **E**.

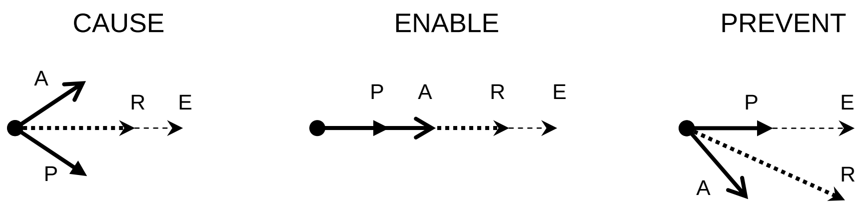

Figure 2. Configurations of forces associated with CAUSE, ENABLE, and PREVENT.

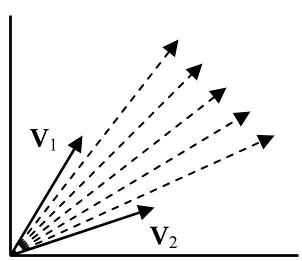

Figure 3. Despite uncertainty about the magnitudes of $\mathbf{V}_1$ and $\mathbf{V}_2$, we can infer that the resultant of the two vectors will reside within the area bounded by $\mathbf{V}_1$ and $\mathbf{V}_2$.

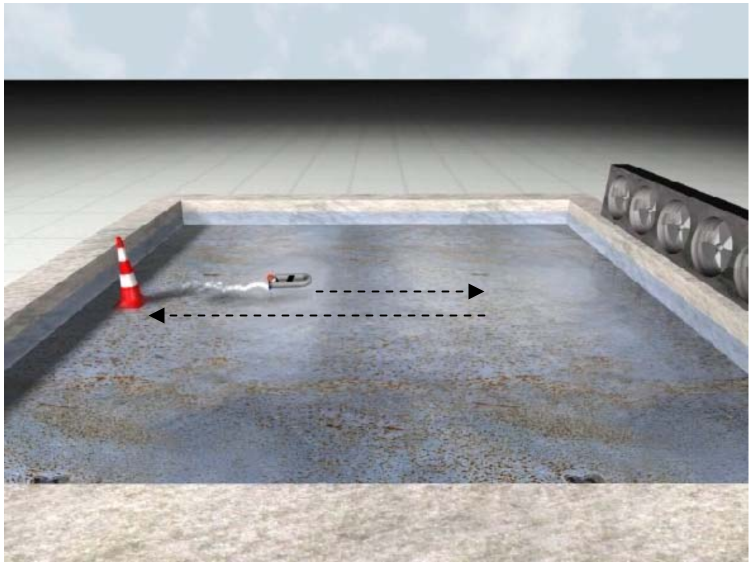

Figure 4: Frame from an animation used to instantiate a CAUSE interaction.

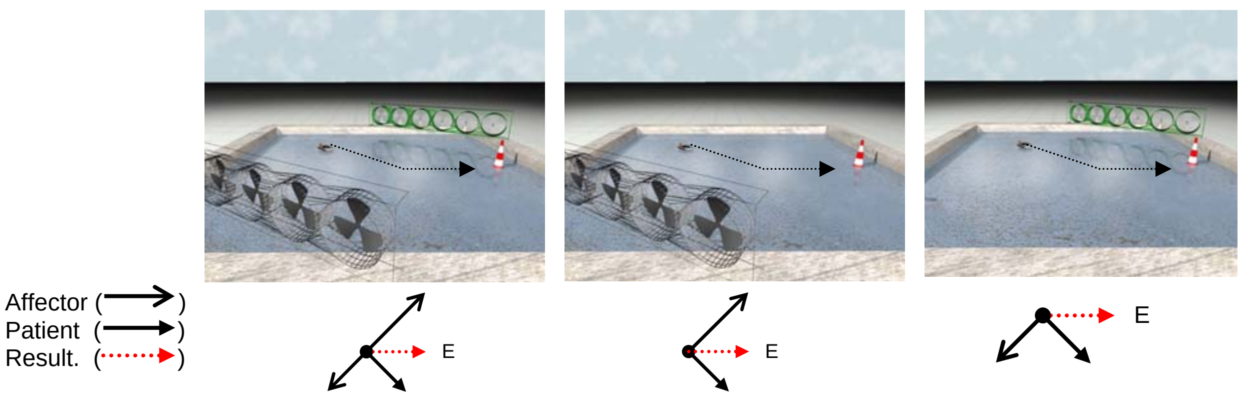

Figure 5. In each animation, the boat motors to the middle, the fans turn on, the boat changes course, and the boat hits the cone. Each animation is based on the same configuration of forces as shown in the first panel. However, in the second and third panels, only one of the two fans appears in the animation, as implied by the incomplete.

CAUSE

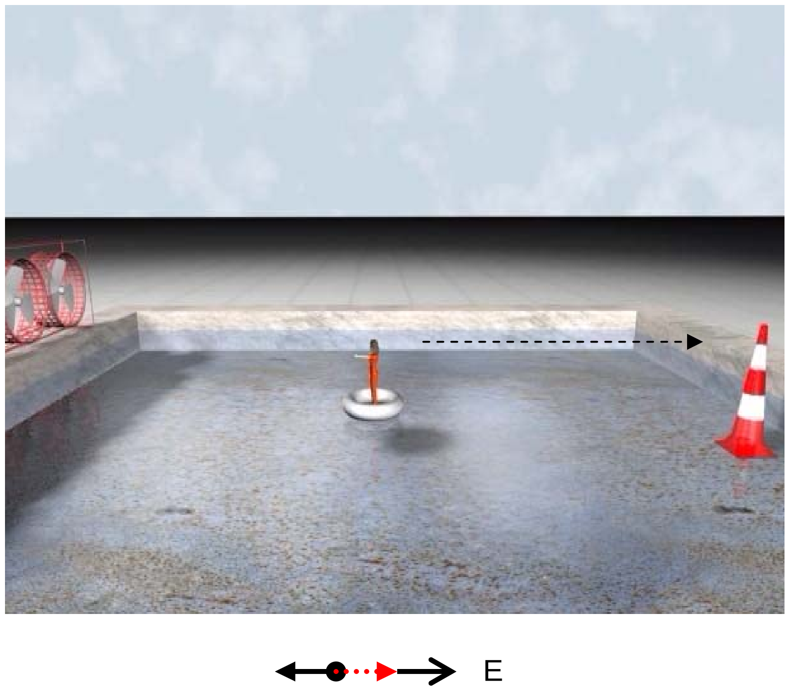

Figure 6. The scene depicts a CAUSE situation since the woman does not want to go to the cone (as indicated by the direction of her pointing) but the fans push her there nevertheless.