Force Dynamics

Phillip Wolff

Robert Thorstad

Emory University

Please address correspondence to:
Phillip Wolff
Emory University
Department of Psychology
36 Eagle Row
Atlanta, GA 30322
Office #: 404-727-7140
Fax #: 404-727-0372
pwolff@emory.edu

In press in *Oxford Handbook of Causal Reasoning* (Michael Waldmann, Ed.)

Abstract

Force dynamics is an approach to knowledge representation that aims to describe how notions of force, resistance, and tendency enter into the representation of certain kinds of words and concepts. As a theory of causation, it specifies how the concept of CAUSE may be grounded in people's representations of force and spatial relations. In this chapter, we review theories of force dynamics that have recently emerged in the linguistic, psychological, and philosophical literatures. In discussing these theories we show how a force dynamic approach to causation is able to account for many of the key phenomena in causal cognition, including the representation of individual causal events, the encoding of causal relations in language, the encoding of causal chains, and causation by omission. Recent accounts of force dynamics have shown how this perspective can be instantiated in computational terms, which not only helps clarify the force dynamic approach to various phenomena, but also helps explain how the units of cognition hypothesized in this approach may be grounded in people's perceptual experience. In addition to reviewing the strengths of this approach, the chapter identifies lines of inquiry for possible future research.

Keywords: Force Dynamics; Causal understanding; Causal Perception; Language and Causation; Causal Reasoning; Causation by omission; Causal composition; Causal chains; Singular causation

Causal events typically involve a large set of factors, many of which may be necessary for the occurrence of an effect (Mill, 1973/1872). However, when people are asked to explain why a particular event occurs, they typically pick out only one or two of the factors (Hart & Honoré, 1959). For example, in the case of a boat capsizing, the captain might point to the occurrence of a rogue wave or a broken rudder, but not the boat's weight, angle to the wind, or the crew. The challenge of identifying which of these conditions constitutes *the* cause of the effect is what Hesslow (1988) has called *the problem of causal selection*.

The phenomenon of causal selection is not just a philosopher's puzzle. It has consequences for our understanding of how people determine guilt in a court case, identify the source of a mechanical failure, or choose a medical treatment. The problem of how people select a particular cause is not easily explained by theories of causation that define causal relations in terms of statistical or counterfactual dependencies (Walsh & Sloman, 2011; Sloman & Lagnado, 2014). To address this problem arguably requires a different kind of theory of causation, one in which causal relations are defined in terms of their internal properties rather than in terms of their external effects. One class of theories of this type are those associated with *force dynamics*. Force dynamics deconstructs the concept of CAUSE and related concepts into finer components. This decomposition not only offers an account of how people solve the problem of causal selection, but also of several other challenging phenomena in the causation literature. We will focus on the representation of causation from a force dynamic perspective, but we will also discuss the force dynamic approach with respect to several other prominent theories of causation. We will argue that force dynamic theories provide an account of the problem of causal selection that is not easily equaled by any other theory of causation.

Force dynamic theories of causation

To explain how force dynamics approaches the problem of causal selection, it will prove useful to review the fundamental assumptions of this perspective through a review of the different accounts of force dynamics.

Talmy's (1988) theory of force dynamics

The first theory of force dynamics to be formulated was proposed by Talmy (1988). The theory introduced the ideas of the imparting of force, resistance to force, overcoming resistance, and removal of a force. As such, force dynamics was hypothesized as a framework that included the notion of causation but also several others notions such as "letting", "hindering", "helping", and "intending." The primary goal in Talmy's theory was to provide an account of the meaning of a large number of verbs, prepositions, and modals. However, Talmy (1988) emphasizes that the patterns of force interactions observed in semantics appear to reflect many of the properties of people's naïve physics. As such, Talmy viewed his theory as not only relevant to the semantics of language, but also as an account of how various causal concepts might be represented in the conceptual system outside of language.

In Talmy's theory, the simplest type of force dynamic pattern is one in which two forces are in steady-state opposition. Such scenarios usually involve two entities. One of these entities, the Agonist, is singled out for focal attention, while the other, the Antagonist, plays a subordinate role. An example of a steady-state pattern would be a tumbleweed that is kept rolling over the ground by a wind. In this scenario, the entity singled out for focal attention would most likely be the tumbleweed, making it the Agonist. The second entity, the wind, would be the Antagonist. While there is movement in such a scenario, there is no qualitative change in the type of action of the Agonist, so the interaction is classified as steady-state. Talmy argued that several kinds of steady-state patterns could be differentiated with respect to three dimensions. One of these dimensions is the *tendency* of the Agonist. The Agonist is associated with a force that gives it a tendency either for rest ~~action~~ or for ~~rest~~action. In Talmy's theory, the two forces in a force dynamic pattern are almost always in opposition; hence, if the Agonist has a tendency for action, the Antagonist is associated with a force that pushes it towards rest, and if the Agonist has a tendency for rest, then the Antagonist is associated with a force pushing it towards action. The second dimension distinguishing different force dynamic patterns is the relative strength of the forces associated with the Agonist and the Antagonist: the Agonist is either stronger or weaker than the Antagonist. The third and final main dimension is the outcome: the Agonist remains either in action or rest. Table 1 summarizes the different kinds of steady-state force dynamic patterns in Talmy's theory.

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

Table 1. Dimensions that differentiate steady-state force dynamic interactions in Talmy (1988)

	Agonist tendency	Agonist strength	Agonist result
'Causative' extended causation	Rest	Weaker	Action
Despite	Rest	Stronger	Rest
Despite/hinder	Action	Stronger	Action
Block (i.e., xtended Prevent)	Action	Weaker	Rest

In steady-state interactions in which the Agonist has a tendency for rest, its strength is weaker than that of the Antagonist, and the result is one of action because the stronger Antagonist force is for action, the overall pattern is expressed in language with predicates implying causation, e.g., *The wind kept the tumbleweed rolling*. In steady-state interactions in which the Agonist has a tendency for rest that is stronger than that of the Antagonist and the resulting outcome is for rest, the over pattern is one of despite, e.g., *The tumbleweed did not move despite the wind*. In steady-state interactions in which the Agonist's tendency is for action and is stronger than that of the Antagonist and the result is one of action, the pattern is a second type of despite, e.g., *The tumbleweed rolled despite the wind*. Finally, in a steady-state interaction in which the Agonist has a tendency for action but is weaker than the Agonist and result is of rest, the pattern is one of prevention, e.g., *The wind prevented the tumbleweed from rolling down the hill*.

A second main type of interaction in Talmy's (1988) theory are change-of-state patterns. In change-of-state patterns, the Antagonist, rather than impinging steadily on the Agonist, enters or leaves this state of impingement. Table 2 summarizes some of the possible change-of-state patterns identified in Talmy's theory.

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

Table 2. Dimensions of change-of-state force dynamic interactions in Talmy (1988)

	Agonist tendency	Agonist strength	Antagonist impingement	Agonist result
Causatives	Rest	Weaker	Into	Action
Preventing	Action	Weaker	Into	Rest
Letting	Action	Weaker	Out of	Action
Letting	Rest	Weaker	Out of	Rest

In change-of-state interactions, when the Agonist's tendency is for rest and the Antagonist's motion is into impingement with the Agonist, resulting in action, the interaction is one of causation, e.g., *The boy knocked the bottle off the wall*. Causation from change-of-state interactions differs from causation in steady-state interactions in instantiating *onset causation*. Another type of change-of-state interaction is one in which the Agonist's tendency is for action, the Antagonist comes into impingement with the Agonist and the result is rest. Such an interaction is one of stopping or prevention, e.g., *The boy prevented the bottle from falling off the wall*. When the Agonist has a tendency for action, the Antagonist moves out of impingement with the Agonist, and action ensues, the interaction is one of letting, e.g., *The boy let the bottle fall off the wall*. Finally, when the Agonist has a tendency for rest and the Antagonist moves out of impingement with it, and the result is rest, the interaction is another type of letting, e.g., *The boy let the bottle sit on the wall*. Whereas onset causation and preventing involve the start of or continuation of impingement of the Antagonist on the agonist, letting involves the cessation of impingement of the Antagonist on the agonist.

Talmy's theory of force dynamics highlights how a relatively large number of verbs and prepositions instantiate various types of force-dynamic patterns. For example, in addition to the verbs *cause, prevent*, and *let*, there are the verbs *keep, refrain, get, stop, make, overcome, push, pull, press, hold, resist, maintain, exert, press, try, hinder, urge,* *persuade, refuse, free, allow, help, permit, forbid, drag, trudge*, and *attract*. Prepositions encoding force dynamics include *against*, *despite*, *although*, *because* and *on*. Talmy's theory explains the ways in which these verbs overlap and differ in meaning. Another key contribution of Talmy's theory is the concept of tendency. It is only via the notion of tendency that the notions of CAUSE, PREVENT and LET can be distinguished: with LET the Agonist's tendency is realized in the final result, whereas in CAUSE and PREVENT it is not. Talmy's theory is conceptually grounded in physical forces and motion events. However, it is also intended to be a theory of social (e.g., *peer pressure, persuade, urge, permit, forbid*) and psychological interactions (e.g., *moral fortitude, overcome, refuse, give in*).

Talmy's theory of force dynamics shows how the concept of CAUSE can be viewed as one member of a family of concepts. However, a closer look at his proposal also raises complications for Talmy's account of the concepts underlying the semantics of various words. In certain ways, the theory invokes more distinctions than are strictly needed in order to differentiate several of the force-dynamic patterns. Consider, for example, the set of dimensions associated with steady-state verbs. Talmy proposes that these interactions are specified in terms of three dimensions. However, knowledge of any two of the dimensions perfectly predicts the value of the remaining third dimension. For example, if the Agonist's tendency is for rest and the Agonist's result is for action, it must be the case that the Agonist's tendency is weaker than that of the Antagonist. A similar kind of redundancy emerges in the case of the change-of-state patterns. As shown in Table 2, in all of Talmy's change-of-state patterns, the Agonist is weaker than the Antagonist. Eliminating this dimension does not change the theory's ability to differentiate the notions of CAUSE, PREVENT, and LET.

In other ways, Talmy's theory appears to lack distinctions that are needed for specifying certain force dynamic patterns or ruling them out. For example, in Talmy's theory, the forces associated with the Agonist and Antagonist are nearly always assumed to be in opposition. If the theory was more accepting of concordance between the Agonist and Antagonist, it might be better able to specify the patterns underlying the notions of *help*, *enable*, and *assist*. Another way in which the theory lacks explicitness lies in its inability to explain why certain force-dynamic patterns do not occur. In the case of the steady-state patterns, the three binary dimensions imply eight possible force-dynamic patterns, but four of these configurations are not discussed. The reason why they are not discussed is presumably because they represent impossible patterns of forces and results. For example, if the Agonist has a tendency for Action and the Agonist is stronger than the Antagonist, then the Agonist's result cannot be one of rest. Such patterns can be recognized as impossible when we attempt to imagine them in our mind. Ideally, their impossibility would be made explicit by the theory. Finally, the various dimensions in Talmy's theory are formulated in terms of the notions of rest and motion. As a consequence, the theory generates redundant versions of the concepts of LET, DESPITE, CAUSE, and PREVENT (one for rest and the other for motion). Such redundancies might be acceptable if evidence could be offered for the existence of multiple versions of these concepts, but no such evidence is offered. As shown in the next theory, these redundancies disappear when the notions of tendency and result are reformulated as relations between forces and an endstate.

Wolff's force theory

In prior work, we have developed a theory of force dynamics that addresses several of the limitations of Talmy's (1988) theory as well as extends his theory to handle several phenomena not considered by Talmy (Wolff & Song 2003; Wolff 2007; Wolff, Barbey, & Hausknecht 2010; Wolff & Barbey, 2015; Wolff & Zettergren, 2002; see also Operskalski & Barbey, this volume). The force theory partitions the representation of causation into two major frameworks: 1) the representation of individual configurations of force and 2) the representation of chains of configurations of forces. Both of these frameworks are important for explaining the main problem addressed in this chapter of how people select the cause from a large set of causal factors.

According to the force theory, individual configurations involve two main entities: a force generator and a force recipient (Wolff, Jeon, & Yu, 2009; Wolff, Jeon, Klettke, & Yu, 2010). We will refer to the force generator as an affector and the force recipient as a patient. The theory proposes that people specify causal relations in terms of configurations of forces in relationship to a vector that specifies the patient's relationship to an endstate. The endstate can be a location in physical space or state in a state space. It is assumed that people's representations of a force specify its source, direction, and relative magnitudes. Absolute magnitudes are not represented and as a consequence each configuration of forces contains a certain degree of uncertainty (Wolff, 2007; 2014). For individual configurations, this uncertainty does not have an impact on the categorization of a configuration of forces because relative differences in magnitude are enough to distinguish different kinds of causal relations. However, when the configurations are combined, it is expected that this indeterminacy can have an impact on how a causal chain is represented, as described below. Lastly, it is assumed that the force in a configuration may be physical, mental (e.g., intentions), or social (e.g., peer pressure) (Copley & Harley, 2015).

Individual relations. At the level of individual configurations of forces, the force theory predicts four main causal concepts: CAUSE, HELP, PREVENT and DESPITE. These three concepts can be differentiated with respect to three dimensions (a) the tendency of the patient for an endstate, (b) the presence or absence of concordance between the affector and the patient, and (c) whether the resultant is directed towards the endstate. Unlike other theories of causation (e.g., Probabilistic theories such as Boddez et al.; Cheng & Lu; Rottman; Rehder), the force theory does not require that the result event occur before it can be said that causation has occurred. The dimension of whether the event actually occurs could be associated with a fourth dimension that is represented in terms of the length of the endstate vector, as described below. Table 3 summarizes how the three main dimensions differentiate the concepts of CAUSE, HELP (also ALLOW and ENABLE), PREVENT and DESPITE (also HINDER). When we say, for example, *High winds caused the man to move toward the bench*, we mean that the patient (the man) had no tendency to move towards the bench (Tendency = No), the affector (the wind) acted against the patient (Concordance = No), and the resultant of the forces acting on the patient was directed toward the result of moving toward the bench (Endstate targeted = Yes).

Table 3. Representations of several concepts in the force theory

	Patient tendency for endstate	Affector-patient concordance	Endstate targeted
CAUSE	No	No	Yes
HELP (also ALLOW and ENABLE)	Yes	Yes	Yes
PREVENT	Yes	No	No
DESPITE / HINDER	Yes	No	Yes

A key feature of the force theory is that it specifies how the dimensions of tendency, concordance, and endstate targeting can be represented in non-linguistic, computational terms. According to the theory, these dimensions specify possible configurations of forces. Example instantiations of these configurations are depicted in Figure 1, which shows scenes involving some wind, a person, and a bench. The configurations of forces are depicted in two ways. In one of these ways, the forces are placed in the scene next to the entity they are associated with: the large open arrow is associated with the affector, the small open arrow is associated with the patient's tendency, and the small solid arrow is associated with the resultant vector of the affector and patient forces. The configuration of forces is also depicted as a free-body diagram positioned below each scene. As is customary in free-body diagrams, the forces are shown acting on only one object, the patient. They do not show the location of the affector, only the direction and magnitude of the affector's force on the patient. In the free-body diagrams, the force associated with the affector is labeled **A**, the force associated with the patient is labeled **P**, the resultant vector is labeled **R**, and the endstate vector is labeled **E**.

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

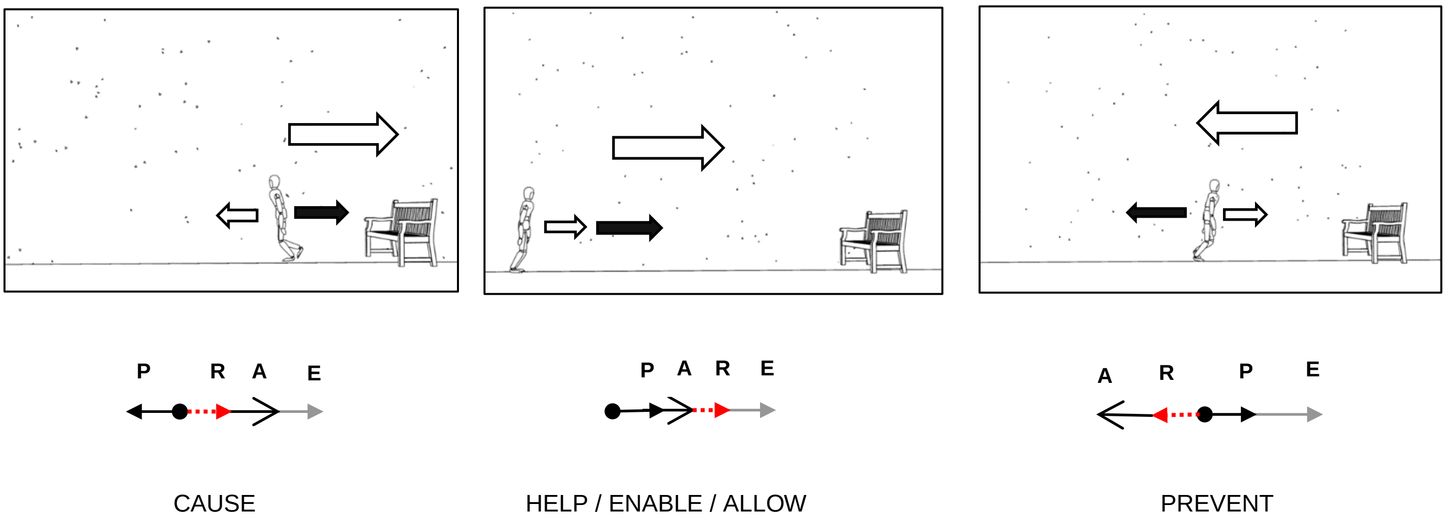

Figure 1. Configurations of forces associated with CAUSE, HELP/ENABLE/ALLOW, and PREVENT; **A** = the affector force, **P** = the patient force, **R** = the resultant force; **E** = endstate vector, which is a position vector, not a force. The images above each configuration depict a scene showing a person, wind, and a bench. In the CAUSE scene, the person's tendency is to move away from the bench, but is pushed back to the bench by the wind. In the HELP scene, the person's tendency is to go to the bench and the window pushes the person in the same direction. In the PREVENT scene, the person's tendency is to go to the bench, but the wind stops the person from reaching it.

The force associated with the patient, **P**, can be generated in a number of ways, including from processes that are internal to patient (e.g., movement of muscles) or from positioning in a force field. Position in a field can give rise to tendencies such as falling due to gravity or to natural processes, such as 'aging', 'ripening' and 'cooling' (Copley & Harley, 2015). The force associated with the patient can also emerge from the patient's resistance to change due to interactions with other entities, as occurs in the case of frictional forces. In Figure 1, the patient's force corresponds to the force generated by chemical potential energy in the patient that allows it to move its muscles. When the patient has a tendency for the endstate, **E**, the patient vector, **P**, points in the same direction as the endstate vector, **E**; otherwise, **P** points in a different direction. When the patient and the affector are in concordance, their respective vectors point in the same direction. Finally, the patient entity will target the end-state when the result (sum) of the **A** and **P** vectors, **R**, is in the same direction as the end-state vector, **E**. The endstate vector, **E**, is a position vector, not a direction vector. Hence, the length of the end-state vector specifies how close the patient is to reaching the endstate. Once the patient reaches the endstate, the magnitude of the endstate vector becomes zero.

The predictions of the force theory have been tested in experiments in which configurations of force associated with CAUSE, ALLOW, PREVENT and DESPITE, among others, were instantiated in 3D animations generated from a physics simulator. As reported in Wolff (Wolff & Zettergren, 2002; Wolff, 2002), people's descriptions of these animations closely matched the model's predictions of how the underlying configurations of force should be classified.

Causal chains in force dynamics. In addition to explaining the representation of individual causal interactions, the force theory also explains how individual relations can be joined to form causal chains and how these chains may then be re-represented as a single overarching causal relation. In the force theory, causal chains are created in one of two ways depending on whether the chain involves transmission or removal of an actual or possible force. In cases of transmission, the resultant force of a configuration serves as the affector force in a subsesquent configuration of forces. The idea can be explained in a simple causal chain in which marble A hits marble B, which hits marble C, as depicted in Figure 2.

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

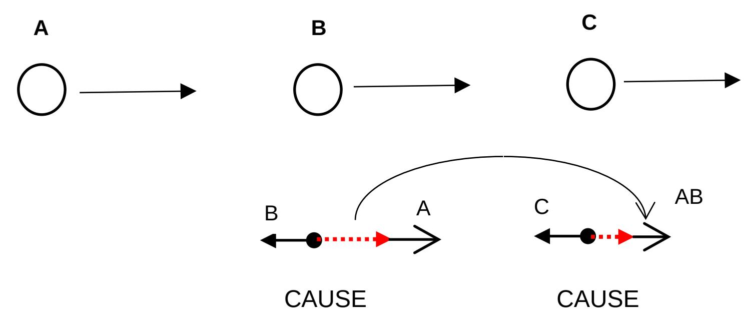

Figure 2. The image above depicts a causal chain of marbles in which marble A causes marble B to hit and move marble C. When A hits B, it results in a CAUSE configuration of forces. The curved arrow shows how the resultant of these forces in the first CAUSE interaction serves as the affector vector in the following CAUSE interaction between marbles B and C. Note, the curved arrow also represents temporal order.

In force dynamic terms, the transmission of force depicted in Figure 2 requires treating the result of the forces associated with marbles A and B as the affector force that acts on marble C.

Whereas some causal chains involve the transmission of force, other causal chains involve the removal of a force (or the nonrealization of a possible force). When a chain involves the removal of a force, the manner in which the resultant forces becomes an affector force reverses from the way it occurs in the ordinary transmission of forces. The removal of a force occurs in situations known as *double preventions*. Consider, for example, a situation in which a force (object or person) knocks out a pole that is holding up a tent so that the tent falls. The pole is preventing the tent from falling, and knocking out the pole prevents this prevention. The ultimate result—the tent falling—is due to the removal of a force, that is, the pole that is keeping the tent up (or preventing it from falling) (Wolff, Barbey, & Hausknecht, 2010; Wolff, & Barbey, 2015).

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

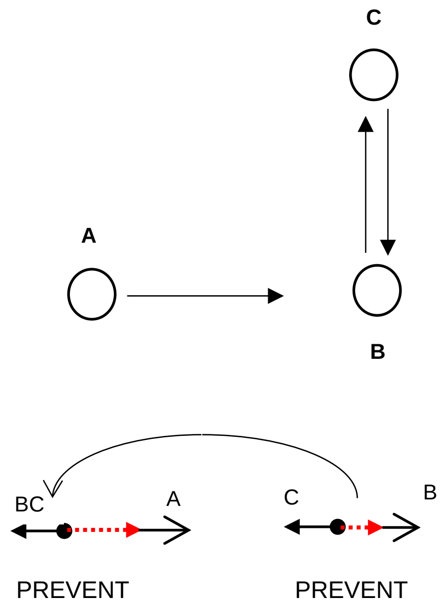

Figure 3. The image above schematically depicts a pole B preventing a tent C from falling. When object A hits B, B can no longer maintain its prevention of the falling, and as a consequence, the tent falls. As implied by the curved arrow, the preventive relationship between B and C must be in place before A can prevent B. In addition, when A acts on B, A acts not just on B, but on the sum of the B and C forces, which in combination form the tendency vector in the interaction of forces that exists between objects A and B.

A chain depicting force removal is shown in Figure 3. In this chain, B represents the pole holding up tent C. When object A knocks out B, B can no longer prevent C from falling, and the tent falls. Notice that when A acts on B, B is already interacting with C. Thus, A acts not only on B, but also on the resultant of the forces associated with B and C. In effect, the resultant of the B and C forces serves as the patient vector force in the interaction between objects A and B.

Double prevention chains are also realized when an affector refrains from applying a force (Wolff, Barbey, & Hausknecht, 2010; Pinker 1989). In such chains, the affector has the ability to prevent the patient from realizing its tendency, but refrains from doing so. In effect, the affector acts on itself to prevent a potential prevention. In Figure 3, refraining can be depicted by renaming the B object A', implying a situation in which object A acts on itself. Imagine, for example, a tent starting to collapse. Someone inside of the tent could entertain the possibility of stopping the collapse by holding up ~~their~~ his arm and serving as a pole. Alternatively, they might decide to let the tent fall down. In choosing the second option, they remove a potential prevention from being realized by holding back the preventive force. In Wolff et al. (2010), such unrealized forces are referred to as *virtual forces*. As shown in Wolff et al. (2010), situations involving virtual forces are described in the same way as situations involving actual forces. People describe situations like the tent scenario with such expressions as *The person allowed the tent to fall* or *The person's inaction let the tent fall*.

A third way in which a double prevention can be realized is through the application of a force that prevents a prevention. Note that this approach to double prevention does not involve the removal of a force nor the holding back from applying a force. Rather, in this kind of double prevention, an affector applies a force on the patient that allows the patient to avoid a preventive force. This third type of double prevention occurs in the cases of bridges. Bridges can apply a (normal) force on a patient (e.g., a car), which allows the patient to avoid a prevention (e.g., a river).

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

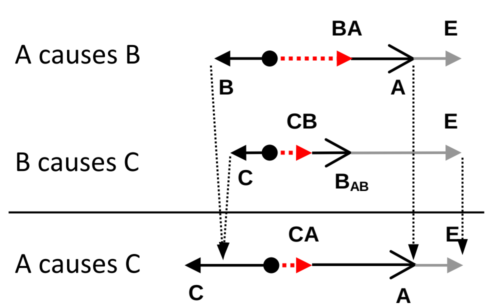

Figure 4. The image above exemplifies how two configurations of forces are combined to form a single configuration of forces. A summary configuration is created using the affector vector from the initial configuration as the affect vector, the vector sum of all of the patient vectors as the patient vector, and the endstate from the last configuration as the endstate vector. Once formed, a resultant vector in the summary configuration is formed by simply summing the patient and affector vectors. The summary configuration as interpreted like any other configuration. In the chain depicted above, a sequence of CAUSE configurations gives rise to another CAUSE configuration.

In extending the force theory to causal chains, it becomes possible to address a number of important phenomena. Central to all of these phenomena is the process by which a causal chain can be re-represented as a single causal relation (see Johnson & Ahn, 2015; Johnson & Ahn, this volume). A simple example of this re-representation process occurs in cases of transitive causal reasoning. When told, for example, that A causes B and B causes C, people can infer that A causes C, or, to use a more contentful example, when told *Water causes rusting* and *Rusting causes discoloration*, people can infer *Water causes discoloration*. Whether the chain involves the transfer or removal of a force, the manner in which summary configurations are derived remains the same: specifically, the affector in the summary configuration is the affector from the first configuration; the end-state is based on the end-state of the last configuration; and the patient vector in the summary configuration is the sum of the patient vectors in the component configurations (see Figure 4). When this procedure is applied to causal chains involving two *causes*, the resulting summary configuration is always a CAUSE relation (Wolff & Barbey, 2015). Interestingly, the process of deriving new causal relations can also occur in chains in which the component relations differ, that is, in chains involving different kinds of causal relations, such as causing, allowing, and preventing. For example, when told that A *causes* B and B *prevents* C, people usually infer A *prevents* C (Goldvarg & Johnson-Laird, 2001; Barbey & Wolff, 2006, 2007; Sloman et al., 2009; Khemlani et al., 2014), or more concretely, when told *Rain causes humidity* and *Humidity prevents evaporation*, people infer that *Rain prevents evaporation*. When causal relations are formed from different kinds of causal relations, the process is not simple transitive reasoning: instead, the reasoning involves a process known as relation composition. Wolff and Barbey (2015) show that the force theory's approach to relation composition is able to predict people's responses to various kinds of relation compositions at least as well as other leading theories of relation composition (Sloman et al., 2009; Goldvarg & Johnson-Laird, 2001).[1]

Causation by omission. The force theory's ability to predict relation compositions allows it to address the problem of how causation can occur from omissions or absences. Causation by omission is causation in which the absence of an influence results in the occurrence of an effect, as in *The lack of light causes depression* or *The absence of water caused the plant to die*. According to a number of researchers, theories like force dynamics cannot account for the phenomenon of causation by omission because such theories require that there be a transmission of energy or force, and clearly no such transfer can occur from an absence (Schaffer, 2000; Schulz, Kushnir, & Gopnik, 2007; Woodward, 2006; Sloman & Lagnado, 2015). In the force theory, however, causation can occur not only when force is transmitted but also when force is removed (or held back). Consider the everyday event of pulling a plug so that water can flow down a drain. Such a situation instantiates a double prevention. First the plug prevents water from draining, which can be viewed as a

---

[1] As document in Johnson and Ahn (2015), not all causal chains are transitive. For example, buying an air conditioner can cause an increase in the electric bill, and an increase in the electric bill can cause anger, but few would want to conclude that buying an air conditioner causes anger. Johnson and Ahn (2015) provide evidence that the transitivity of a chain depends on the ability to fit the relations into a single schema, which may exist by virtue of a common underlying mechanism, temporal contiguity, and homogeneity of the time scale of the different relations.

pre-existing condition (Cheng & Novick, 1991). Next, an agent removes the plug, thereby removing (i.e., preventing) the prevention of the water from flowing down the drain. Importantly in cases of double prevention such as this, the situation can always be described in terms of absences. In the case of water flowing down a drain, it can be said *The absence of the plug caused/allowed the water to flow down the drain*. Such expressions of causation by omission are possible because double preventions are realized through the removal of a force, which creates an absence. In a series of studies in Wolff et al. (2010), we found evidence for this proposal: when shown animations of double preventions, people endorsed statements stating that the *lack* of the second entity in a double prevention (e.g., the plug) allowed or caused the third entity (e.g., water) to undergo a process leading to the result.

Allow relations. Following McGrath (2003), we propose that ALLOW relations are based on double preventions (Wolff et al., 2010; Wolff & Barbey, 2015). In the simplest case, ALLOW relations involve removing a force that was originally preventing an event from happening. In the water and plug example, we can say *The removal of the plug **allowed** the water to drain*. The notion of ALLOW is related to the notion of HELP. As discussed earlier, a HELP configuration is one in which the affector force is concordant with the patient force, as implied in sentences such as *The Boy Scout helped the grandmother cross the road*. ALLOW relations also imply concordance between the affector and the patient when analyzed with respect to the resulting summary configuration. The underlying prevent relations in ALLOW entail that the affector's influence is necessary for the occurrence of the effect. The affector's influence is necessary because in a double prevention, the occurrence of the final result is in some way blocked and removal of that blockage depends on the affector. This account of ALLOW was supported in a set of studies described in Wolff et al. (2010). Participants viewed animations instantiating double preventions and, as predicted, they endorsed statements asserting that the first entity in the double prevention allowed the last entity in the double prevention to undergo a certain result. In Wolff et al. (2010), it was also shown the semantics of the verb *enable* is much the same as those of *allow*, suggesting that there may be a set of verbs based on double prevention, including the verbs *allow, enable, let*, and *permit*.

As discussed earlier, double prevention can be instantiated in multiple ways. In addition to the removal of a force, a double prevention can be instantiated by an affector refraining from applying a preventive force (e.g., standing aside to let someone pass) or by the application of a force that prevents a preventative force from being experienced by the patient (e.g., a bridge exerting a force on a traveler so that they do not experience the preventative force of, for example, the river). Regardless of how the double prevention is realized, the relationship can still be viewed as one of *allowing*, *enabling*, or *letting*.

Relationship between Talmy's force dynamics and Wolff's force theory. The force theory differs from Talmy's account of force dynamics in several ways. First, the force theory uses the dimension of concordance between the affector and patient, whereas Talmy's theory generally restricts force interactions to those in which the affector and patient forces oppose one another. Whereas the force theory postulates that the concept of HELP involves concordance between the affector and the patient, Talmy's force dynamics holds that the concept of HELP involves extended disengagement of the Antagonist. As such, Talmy's theory seems to miss the sense of agreement or concordance between the affector and patient in the meaning of verbs like *help*. In Talmy's theory, the concept of LET involves the removal of a force via the removal of the Antagonist. In the force theory, the affector (or Antagonist) removes some other entity from the situation (sometimes the affector itself), thereby removing a force from the patient (or Agonist). According to the force theory, LET necessarily involves at least two configurations of forces, and hence involves a causal chain. In Talmy's theory, letting does not require a causal chain, but rather involves a change in the nature of the interaction between the Antagonist and Agonist over time (i.e., a change-in-state). According to the force theory, the CAUSE configurations may be simpler and more direct than LET configurations, whereas in the Talmy's theory, CAUSE events are no less complex or indirect than LET events. Third, the force theory postulates a set of dimensions that explains how the notions of CAUSE, HELP, and PREVENT are related to one another. In Talmy's force dynamics, these notions are represented by different sets of dimensions, obscuring their similarities and differences. Third, Talmy defines the dimensions of tendency and result with respect to rest and motion, which leads to multiple versions of several causal concepts (one for rest and the other for motion). In the force theory, tendency and result are defined with respect to a location in space, and as a consequence, redundant versions of various causal concepts do not arise. Fourth, the force theory is explicit enough to be computational, and as a consequence, it is able to explain why certain parameterizations of the dimensions are never realized: certain configurations of forces are never realized because they are ruled out by vector addition. The force theory also allows for the instantiation of the various configurations of force in a physics simulator, hence showing how different causal concepts are grounded in the physical world (see Wolff, 2007). Lastly, the force theory makes explicit how individual configurations of forces can be combined to form causal chains, which allows the theory to explain relation composition, causation by omission and the concept of ALLOW.

Copley and Harley's (2015) force-theoretic model

In linguistics, it has been hypothesized that the concept of CAUSE might be essential in the argument-structure properties of various classes of verbs (Van Valin, 1990; Hale &

Keyser, 1993; Marantz, 1997). Consider, for example, the expression *Peter melted the butter*. The expression can be paraphrased *Peter caused the butter to melt*, implying that the expression encodes the notion of CAUSE. Such Vendlerian accomplishments are usually analyzed as composed of two subevents connected by a causal relation. The subevents include a causing subevent e_1, Peter's actions, and a result subevent e_2, the melting of the butter (Pustejovsky, 1995; Dowty, 1979). In formal semantics, such expressions can be represented as $\exists e_1 \exists e_2$: e_1 CAUSE e_2. Such formal representations work well in the case of English because in English accomplishments entail the occurrence of the final result (or event). However, in many other languages, accomplishments do not entail the final result, giving rise to the phenomenon of non-culminating accomplishments. For example, in Karachay-Balkar, a Turkic language spoken in Russia, it is possible to say *Kerim opened the door, but he did not succeed*, and in the Salish language St'a´t'imcets it is possible to say *I made the basket, but it didn't get finished* (Copley & Harley, 2015). A related phenomenon occurs in English with the verbs *enable*, *allow*, and *let*. To say that *John allowed Susan to cut the lawn* implies that the final result probably occurred, but it is not strictly entailed, as evidenced by the acceptability of the sentence *John allowed Susan to cut the lawn, but the lawn wasn't cut*. In Jackendoff's (1991) system of force dynamics, verbs such as *allow* specify that the outcome is *undetermined*. In languages other than English, an undetermined outcome holds not only for verbs like *allow*, but also for verbs like *cause*. Non-culminating accomplishments appear to be quite common across the world's languages (Copley & Harley, 2015; Copley & Wolff, 2015). Assuming the notion of causation is handled in similar ways across languages, the occurrence of non-culminating accomplishments suggests that the traditional approach to how the argument structure of accomplishments is represented might need to be changed.

As discussed earlier, in the force theory, a CAUSE configuration of forces can exist without the occurrence of a result. Copley and Harley (2015) use this property of forces to develop a new theory of argument-structure representation that is able to handle the phenomenon of non-culminating accomplishments. The key ideas in their proposal are the replacement of events with situations and the use of forces as functions that connect an initial situation to a final situation, assuming the absence of external interventions. In Copley and Harley's (2015) force-theoretic model, forces have their origins in all of the individuals and properties in a situation. As such, Copley and Harley's (2015) theory is well suited for explaining causal relationships involving ambient causes and effects, such as *Low interests rates cause inflation*. In the force-theoretic model, causal chains are represented through the repeated application of net forces on different situations, as depicted in Figure 5. One major difference between the force-theoretic model and Talmy's model is that in the force-theoretic model, the notion of tendency is not represented. While the use of a net force allows the model to handle more abstract kinds of causation, the absence of the notion of tendency means that the model is unable to distinguish several kinds of causal notions, such as the difference between CAUSE and HELP.

While the force-theoretic model fails to make certain distinctions, its generality allows it to explain the argument-structure properties of a number of verb categories. In addition, the model might serve well as an account of how force dynamic representations emerge developmentally in children. In particular, Göksun, George, Hirsh-Pasek, & Golinkoff (2013) found that while 3.5 – 4.5-year-olds were relatively good at representing configurations of force associated with CAUSE, but they struggled with ENABLE and PREVENT. Göksun et al.'s experimental design allowed them to determine that young preschoolers have difficulty representing multiple forces. It was not until 5.5 years-of-age that children were able to represent the full range of two-force configurations associated with CAUSE, ENABLE, and PREVENT. Göksun et al.'s findings are consistent with the possibility that force dynamics in young children may not represent the tendency of the patient and instead might involve just one force. In support of this phenomenon, Bowerman (1982) found that children often confuse the verbs associated with causing and letting.

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

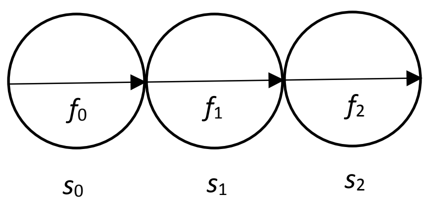

Figure 5. The force-theoretic model represents causal chains composed of a sequence of net forces (f_0, f_1, f_2) that bring about a sequence of situations (s_0, s_1, s_2).

Gärdenfors (2014) Two-vector model

Verbs specify only certain aspects of an event or state (Pinker, 1989; Wolff, 2003, 2012; Wolff & Malt, 2010; Jackendoff, 1990), and this selectivity often forms identifiable patterns. One pattern that has been observed is that verb meanings appear to specify the notion of either manner or result, but not both, a phenomenon referred to as manner/result complementarity (Levin & Rappaport Hovav, 2011; Rappaport Hovav & Levin, 2010). Consider, for example, the event depicted in Figure 6 in which a suitcase is pulled over a line by a person.

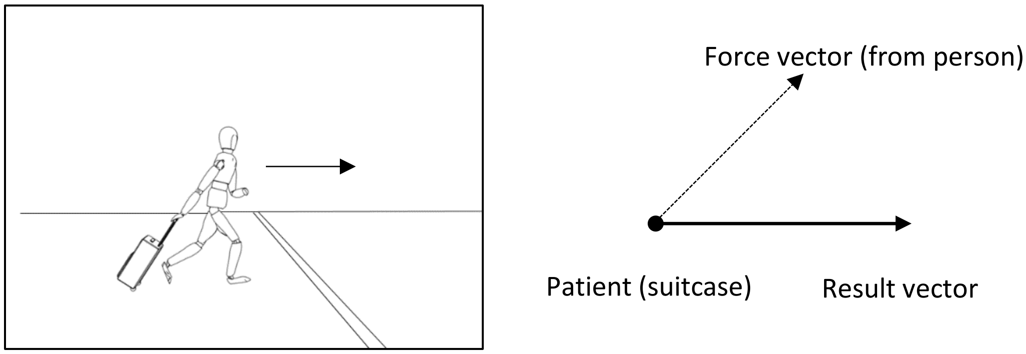

Figure 6. In this scene, a suitcase is pulled over a line.

It could be said *The suitcase crossed the line* or *The suitcase rolled*. Interestingly, there is no one verb in English that codes for both crossing and rolling. Gärdenfors (2014; see also Warglien, Gärdenfors, & Westera, 2012) have developed a force-dynamic account of why manner and path components do not appear together in the meaning of a verb. According to their *two-vector model*, verb meanings are based on vectors. However, a verb can encode multiple vectors only if they are from the same domain, a restriction they refer to as the single-domain constraint (Gärdenfors, 2014; Warglien et al., 2012). Gärdenfors defines a domain as a set of dimensions that are integral rather than separable. Dimensions are integral when having a value on one dimension entails a value on the other dimension (Nelson, 1993). For example, pitch and loudness are integral because if a sound has a pitch, it must also have loudness. Pitch and hue are not integral because a value on one does not entail a value on the other. Turning to force and result vectors, such vectors come from different domains because they are based on separable dimensions. Thus, according to the two-vector model, verbs encode either a result vector or a force vector, but not both.

There are several similarities between the two-vector model and the force theory (Wolff, 2012). In particular, both models propose that verb meanings are based on vectors and both propose force vectors and result vectors. There are also several ways in which these models differ. In the two-vector model, the result vector indicates whether a result occurs, whereas in the force theory, the result vector (or resultant vector) indicates whether the patient will move towards the endstate, but not whether the patient ultimately reaches the endstate. In the force theory, whether the patient reaches the endstate is coded in the length of the endstate vector. A second way the two models differ concerns the kind of phenomena the two models seek to explain. The two-vector model offers an account of the difference between manner and result verbs and of manner/result complementarity. The force theory, in contrast, offers an account of the meaning of a particular class of verbs and prepositions associated with the expression of causation and related notions. Given these differences in emphasis, the models can be viewed as at least partially complementary. However, there is at least one way in which the assumptions of the two-vector model conflict with those of the force theory.

The two-vector model rules out combinations of vectors from different domains. This property of the two-vector model would appear to rule out the kinds of representations hypothesized in the force theory. According to the force theory, verb meanings are based on configurations of vectors that combine force vectors with position vectors (for coding an endstate). The two-vector model rules out such configurations because force and position vectors come from different domains.

We would argue that the single-domain constraint is too restrictive. Without the ability to relate force vectors to position vectors, it is not possible to express differences in meaning between various kinds of causal verbs, such as CAUSE and PREVENT, which, according to the force theory, differ only in the direction of their endstate vector. Such a constraint also rules out the ability to distinguish verbs like *push* and *pull*. Clearly, the verbs *push* and *pull* code for force. According to Zwarts' (2010) vector analysis of these verbs, the information that differentiates the meaning of these verbs is in the direction of the force vectors *relative* to a spatial relation between the agent and the patient. In both *push* and *pull*, the agent's force on the suitcase is exactly the same. What differs is that in pushing, the force is directed away from the agent, whereas in pulling the force is directed towards the agent (Zwarts, 2010). In the force theory, this critical difference between the verbs would be specified in the endstate vector (Wolff, 2012). Zwarts' analysis of several other verb pairs (e.g., *squeeze* and *stretch*, *lean* and *hang*) suggests that the meanings of these verbs are based not only on forces, but also on position vectors that specify the spatial direction of those forces. While the single-domain constraint appears to be too restrictive, it may be possible to reformulate the constraint such that it is possible to explain manner/result complementarity without undermining the ability to represent the meanings of various force dynamic verbs. According to the force theory and Talmy's theory of force dynamics, all of the forces are understood with respect to the patient: the patient serves as the origin for all of the forces in the interaction. Manner verbs, in contrast, typically specify forces acting on the affector. Manner/result complementarity might result from a constraint in which verbs can only specify forces with respect to a single actor in a situation, either the affector or the patient.

Mumford and Anjum's (2011) Dispositional theory of causation

One of the key concepts in force dynamics is the notion of tendency. A tendency can be thought of as a property of an object that grants it a disposition. Tendencies or dispositions factor into the realization of an effect. The basic idea has its origins in Aristotle's causal powers approach to causation (Wolff & Shepard, 2013). Aristotle emphasized that in causal interactions, both the agent and patient had causal powers. The agent had the ability to transmit a causal power and the patient, the capacity to receive a change (Marmodoro, 2007). A recent version of this approach to causation is reflected in Mumford and Anjum's (2011; Anjum & Mumford, 2010; see also Harré and Madden, 1975) theory of *causal dispositionalism* (for review see Waldman & Mayrhofer, 2016).

In Mumford and Anjum's theory, objects have causal powers by virtue of their properties. For example, putting a penny on a scale causes the needle of the scale to change. It changes because the penny contains the property of weight or mass. It also possess several other properties like shape and color, but these properties have no causal impact on the scale. Adding sriracha sauce to one's rice noodles causes them to taste hot. The chili peppers and salt in sriracha give it the property of spiciness and it is this property that brings about a certain taste, not the sauce's red color or thick consistency. Objects roll due to their spherical shape and stay put due to their flat sides. Some objects have the power to hold liquids due to their concave shapes. According to Mumford and Anjum (2011), objects have dispositions that manifest themselves in properties, like fragility, weight, warmth, smoothness and momentum.

In Mumford and Anjum's (2009, 2011), theory, just as in Aristotle's theory, both agents, like the penny, and patients, like the scale, are endowed with causal powers. Mumford and Anjum suggest that causal powers can be modeled as vectors. Just like vectors, causal powers have a direction and a magnitude. In addition, vectors can be added together to give rise to a resultant vector. An example of Mumford and Anjum's approach is shown in Figure 7. The starting point for the situation is a vertical line drawn in what is referred to as a quality space. When the resultant vector reaches a threshold, a result occurs.

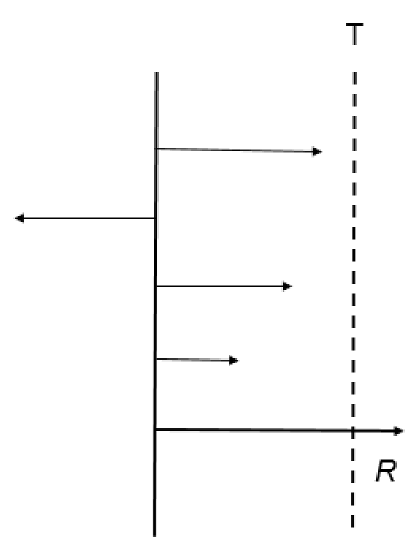

Figure 7. The addition of different causal powers (e.g., powers associated with lightning, rain, oxygen, and matches) combine to form a resultant causal power with a great enough magnitude to surpass a threshold triggering an effect (e.g., fire).

Mumford and Anjum's (2009, 2011) causal disposition theory of causation differs in several ways from the theories based on force dynamics. First, the vectors in Mumford and Anjum's theory are causal powers rather than forces. Second, the theory does not differentiate between agents and patients. The cause of a ball's rolling could just as well be due to its spherical shape as to the cue stick hitting it: all such objects in the situation have causal powers that can be added together to give rise to a resultant causal power. Another difference is that effects are triggered only when the resultant reaches a threshold. The notion of a threshold is not present in theories based on force dynamics.

Despite these differences, there are also several commonalities between Mumford and Anjum's theory and theories based on force dynamics. The most significant commonality is the idea that causal interactions depend on tendencies, or in Mumford and Anjum's theory, dispositions. A second commonality is the use of vectors in the modeling of causal influences. In Mumford and Anjum's theory, the vectors represent causal powers, which cover a wide range of causal influences than forces, but clearly, the way the notion of force is used in several theories of force dynamics overlaps with the notion of causal power as used in Mumford and Anjum's theory. Finally, theories of force dynamics and the causal disposition theory both hold that effects follow from the summation of causal influences in a situation.

While there are several commonalities between force dynamics and Mumford & Anjum's account, force dynamics makes the ontological commitment that the causal powers involved in causal relations are forces. One line of evidence for this claim comes from perceptual priming experiments. Wolff and his colleagues (Wolff & Shepard, 2013; Wolff, Ritter, & Holmes, 2014) showed participants animations involving either causal interactions, e.g. a collision of two marbles, or non-causal interactions. While viewing the animations, participants held a haptic controller that sometimes created a small force against the participant's hand. The authors found that participants detected physical forces against their hand more rapidly when viewing causal than when viewing non-causal animations, i.e. that viewing causes primes detecting physical forces. This finding lends support to the claim that causal powers may be represented specifically in terms of forces.

While Mumford and Anjum's disposition theory is well suited for capturing intuitions about the complexity of causal interactions, with all of the factors that enter into them. Mumford and Anjum also argue that their theory is able handle probabilistic causation, causal chains, various phenomena associated with the perception of causation (Mumford and Anjum, 2011). However, as explained in Mumford and Anjum (2011), the theory has difficult with causation by omission. In addition, since it does not differentiate causal roles, like affector and patient, the theory is unable to distinguish different kinds of causal relations. It arguably makes up for these limitations in its detailed account of the notion of dispositions.

Pinker's theory of force dynamic relations (1989)

In Talmy's 1988 theory of force dynamics, the distinction between steady-state or extended causation (e.g., *The wind kept the tumbleweed rolling*) and "onset causation" (e.g., *The wind caused the tumbleweed to start rolling*) is emphasized. In Wolff's force theory, it is not. Indeed, in Wolff's force theory there is no explicit device for indicating whether a change in state or location has occurred, only whether the resultant vector has targeted an endstate. It certainly would be possible to augment the force theory with the ability to specify changes in state or location. In particular, a telicity dimension could be added to the model by use of the endstate vector. If the endstate vector changes length, it could be said a change has occurred, and if the endstate vector changes to a length of zero, it could be said that a particular endstate was reached. Such a vector would be much the same as the result vector proposed in Gärdenfors (2014) two-vector model. This dimension was not added to the force theory because the assumption was that various concepts of CAUSE do not entail a change in location or state. Indeed, it is because the force theory allows for causation without the reaching of an endstate that it is able to handle the phenomenon of non-culminating accomplishments, as discussed in reference to Copley and Harley's *force-theoretic model*. Steady-state causation illustrates a closely related phenomenon. Consider the situations described in the sentences in (1).

1a. Air particles cause a balloon to remain inflated.

b. Fear caused them remain motionless.

c. Flooding caused the museums to stay closed.

d. Small ridges cause water to stand on the concrete.

e. Keels cause sailboats to stay upright.

In each of the situations described in (1), nothing happens. There is no regular sequence of events, overt transfer of conserved quantities, or change in state, and yet the situations can still be construed of as causal. What is true of each of these situations is that they instantiate a configuration of forces. From a force dynamic point of view, it is this configuration of forces that makes them causal, even in the absence of any change.

The ability to represent steady-state situations is very nicely explained in Pinker's (1989) force dynamic account of causal relations. Pinker's (1989) inventory of possible causal links is shown in Table 4. Pinker proposes that the different kinds of causal links are decomposed with respect to four features. The feature of *focus* concerns whether the link emphasizes the cause or the effect in the causal relation. Causal relationships that emphasize the effect are realized in expressions using subordinating conjunctions, including *because, despite, after*, and *when* (Wolff, Klettke, Ventura, & Song, 2005). In the sentences in (2), for example, the main clause expresses an effect and the subordinated clause expresses the cause. In other expressions of causation, as referred to by the mnemonics 'effect', 'but', 'let', and 'prevent' in Table 4, the effect is subordinated to the cause.

2 a. Jerry missed his flight because the taxi got lost.

b. He died soon after, despite receiving the best possible care.

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

Table 4. Inventory of different types of causal links (Pinker, 1989)

	Mnemonics for types of links	Focus	Potency	Cause occurrence	Effect occurrence
Onset					
	'effect'	Cause	Success	Yes	Yes
	'because'	Effect	Success	Yes	Yes
	'despite'	Effect	Failure	Yes	No
	'but'	Cause	Failure	Yes	No
	'let'	Cause	Success	No	Yes
Steady-state					
	'prevent'	Cause	Success	Yes	No

A second dimension proposed by Pinker is *potency*. Potency is a *success* when the antagonist succeeds in exerting its effect over the agonist, otherwise it is a *failure*. Examples in which potency is a failure include sentences using the preposition *despite*, as in *Ralph bought the T.V. despite the wishes of his wife*. In this scenario, the affector, his wife, is not successful in blocking the patient, Ralph, from buying a T.V. Another type of scenario where potency is a failure occurs in sentences containing the conjunction *but*, which imply the notion of trying without success, as in *Terry pushed against the truck, but was unable to make it move*. In this scenario, the affector, Terry, exerts force on the patient, the truck, but the desired outcome does not occur. The key difference between 'despite' and 'but' scenarios is that in ‘despite’ scenarios, the focus is on the effect whereas in ‘but’ scenarios, the focus is on the cause, as is reflected in whether the cause or effect appear as the subject of the sentence. In Pinker’s (1989) proposal, the common notion underlying the verbs *enable*, *permit*, and *allow* is the idea of the affector ceasing to engage with the patient, and as a consequence, something happening to the patient. Pinker refers to this category of causal relations with the mnemonic ‘let’. In ‘let’ scenarios, the cause occurrence dimension is ‘no’, while in in the remaining scenarios, the cause occurrence dimension is ‘yes’. Prototypical causation is reflected in situations in which the focus is on the cause, potency is successful, the affector engages with the patient, and the final effect occurs. Pinker (1989) refers to this kind of relation with the mnemonic ‘effect’. In all of the interactions described so far, the type of causation has been onset causation. Pinker’s model expresses steady-state causation by focusing on situations in which the focus is on the cause and the potency is successful, but it is successful in the sense that a final effect does not occur. Such situations are reflected in the definition of verbs like *support*, *keep*, *suspend* and *occupy*. Such situations describe an affector continuously exerting a force on an agonist such that a particular effect, like falling, is prevented from occurring. It is for this reason that Pinker refers to these steady-state situations with the mnemonic ‘prevent.’

Pinker’s (1989) account of force dynamic relations overlaps with the proposals of Talmy and Wolff. In particular, all three accounts code for the occurrence of the effect, although in Wolff’s force theory, the occurrence of the effect is only implied rather than strictly entailed. Both Talmy’s and Pinker’s accounts distinguish steady-state from onset causation, whereas force theory does not. Talmy’s theory allows for several different kinds of steady-state scenarios, whereas Pinker’s account specifies only one kind of steady-state. In the force theory, letting scenarios are necessarily more complex than the other force dynamic relations because they necessarily involve at least two stages. This difference in complexity is not explicitly captured in Pinker’s feature set. Pinker’s (1989) potency dimension is unique to his proposal. Interestingly, it appears that this dimension can be captured Wolff’s force theory by coding for whether the affector force is concordant with the resultant force. Pinker’s theory goes further than Talmy’s and the force theory in its specification of the notion of trying. However, as described above, it may be possible to capture the notion in the force theory with the endstate vector. One problem with Pinker’s theory is that it does not seem to have a representation for an onset version of PREVENT. If the representation for ‘prevent’ is used to specify onset PREVENT, Pinker’s theory is no longer able to distinguish between steady-state and onset causation. Pinker’s focus dimension is also unique to his proposal. In Talmy’s account and the force theory, the focus is always on the patient (or agonist). The focus dimension allows Pinker’s theory to account for a difference we see in the syntax of different causal expressions, namely whether the cause is subordinated to the effect or vice versa.

Commonality between the different theories of force dynamics

The theories of force dynamics differ markedly, but they also share some deep commonalities. Most of these theories distinguish a relatively wide range of causal concepts, with the concept of CAUSE being just one member of a family of concepts. A second deep commonality is that they all have built into them the roles of force creators and force recipients (see Beavers, 2011; Rappaport Hovav & Levin, 2001). Finally, all of the theories of force dynamics offer at least an intuitive account of the notion of mechanism. Each of these commonalities appears to be essential to explaining how people identify the one factor in a situation that is identified as *the* cause of an effect.

Causal selection problem

From a force dynamic perspective, determining the cause of an event depends on at least three properties of that event.

Causes versus enabling conditions. The first is the nature of the relationship between a candidate causal factor and the effect. In order for a factor to be the primary cause of an event, it must be a cause rather than enabling condition. Consider, for example, a situation in which a person drops a match in dry grass the result is a forest fire. One of the factors in this situation would be the match, another the oxygen in the air. Both the match and oxygen are necessary conditions for the fire, because without either, the fire would not occur. However, the match is more readily construed as the cause of the forest fire and oxygen is more easily construed as an enabling condition. As many have noted, the difference between the match and oxygen cannot be explained in terms of necessity or sufficiency (Cheng & Novick, 1991, 1992; Einhorn & Hogarth, 1986; Goldvarg & Johnson-Laird, 2001). As discussed earlier, force dynamic theories provide an account of the difference between causes and enabling conditions. According to the force theory in particular, enabling conditions are based on double preventions. As already discussed, there are several ways in which double preventions can be realized: they can occur from the removal of a force, the refraining of application of a force, and the application of a force that prevents a prevention from being realized. In the case of oxygen, the double prevention seems to be based on the application of a force that prevents a prevention from being realized. If oxygen were not present, the match would be prevented from igniting.

There have been other accounts of the difference between causes and enabling conditions. For example, in Cheng and Novick's (1991, 1992) probabilistic contrast model, causal relation are based on covariation observed within a "focal set" of events (rather than the universal set of events). Causal relations are associated with positive covariation as indicated by the probability of the effect in the presence of the candidate cause, $P(E|C)$, being noticeably greater than the probability of the absence of the cause, $P(E|\neg C)$. In contrast, an enable enabling condition is inferred for causal factors that are constantly present in the reasoner's focal set (making $P(E|\neg C)$ undefined), but covary positively with the effect in another focal set. Oxygen, for example, is constantly present in most situations, but if it were varied, the effect of fire would vary with its presence. From the point of view of force dynamics, Cheng and Novick's account of enabling conditions makes sense. As discussed above, many, if not most, double preventions have pre-conditions: a preventive force needs to be present before it can be removed. Force dynamic theories therefore predict (albeit weakly) that enabling conditions will tend to have pre-existing conditions. Because pre-existing conditions can last, they may be perceived as constancies in a situation, and may therefore may accord with the main claim of Cheng and Novick's account of enabling conditions.

Force creators. One property of causal situations that is uniquely predicted by all force dynamic theories of causation is that the primary cause of an event must be a force creator (Wolff, Jeon, Klette, & Li, 2010; Wolff, Jeon, & Li, 2009). Consider, for example, the sentences in (3).

3 a. The chef smashed the hot potato.

b. The knife cut the hot potato.

c. The fork lifted the hot potato.

In English, both (3a) and (3b) are acceptable descriptions of causal events, but (3c) is not. From the point of view of many current theories of causation, this difference in acceptability is hard to explain. For example, as discussed in reference to Cheng and Novick's (1991, 1992) probabilistic contrast model, a cause is a factor that increases the likelihood of an effect, that is, that makes $P(E|C)$ greater than $P(E|\neg C)$. Arguably all of the affectors used in (3) increase the likelihood of the effects referred to in the sentences. The unacceptability of (3c) cannot be explained as due to the fact that the named affector is inanimate. Sentence (3b) shows that inanimate entities can sound perfectly fine as causers, at least in languages like English (Wolff, Jeon, Klette, & Li, 2010; Wolff, Jeon, & Li, 2009). In addition, the unacceptability of (3c) cannot be explained as due to the affector being causally unrelated to the effect. As shown in the sentences in (4), both knives and forks can serve as instruments in sentences describing causal chains.

4 a. The cook cut the hot potato with a knife.

b. The cook lifted the hot potato with a fork.

According to Wolff et al. (2010), the reason why the sentences in (3a) and (3b) sound acceptable while the sentence in (3c) is not acceptable is because the affectors in (3a) and (3b) can be construed of as force creators, whereas the affector in (3c) cannot. In the sentence in 3b, the knife creates a force by pushing the parts of the potato apart, while in 3C the force is really not created by the fork but by the person controlling the fork. The sentences exemplify the hypothesis that in order for a causal factor to be the primary cause of a situation, it must be a force creator.

There are several ways in which forces can be created (Wolff, Jeon, Klette, & Li, 2010). First, forces can be created through energy conversion, that is, when energy is transformed from one form to another (Young & Freedman, 1999). For example, the forces involved in lifting one's hand into the air or pushing it down to smash a potato begin with a transformation of potential energy, in the form of chemical potential energy, into motion, kinetic energy. Energy transformation can also occur in inanimate entities. In internal combustion engines, energy conversion occurs when chemical potential energy in gasoline is transformed into kinetic energy. Entities that generate create forces from energy transformation—intentional agents, natural forces, power devices—seem to make good causers as indicated by the finding that they are almost always construable as causers across a wide range of languages (Wolff et al., 2010). Mayrhofer and Waldmann (2014) found empirical support for the hypothesis that the entities that make good causers are those that can generate their own force. In their experiment, participants saw collision events involving two balls. However, prior to the collision, the balls moved in ways intended to imply different degrees of agentivity (e.g., self-propelled motion). Mayrhofer and Waldmann found that increasing the first object's agentivity resulted in participants being more willing to describe it as causing the motion of the second object. White (2006) obtained a similar finding in discovering that participants were more willing to say that an object X caused an object Y to move when object X moved before object Y (but see Hubbard & Ruppel, 2013). The importance of the affector being a force generator is implied in studies examining the role of intention and causality (see Wolff, 2003). The prototype of a force generator is most likely an intentional entity capable initiating its own forces. In support of this hypothesis, Muentener and Lakusta (2011) found that children described more events as causal when they were (A) intended versus (B) unintended or (C) caused by an inanimate object. Muentener and Lakusta (2011) propose that in producing and comprehending language, children may have an intention-to-CAUSE bias.[2] In sum, evidence from multiple sources converges on the conclusion that entities that are able to create forces through energy conversion are readily viewed as potential causes of an event. However, energy conversion is not the only way in which forces can be generated.

A second way in which a force can be created is through physical contact. When an object hits another object, it imparts a force. Crucially, the imparted force does not exist

[2] Muentener and Lakusta's (2011) Experiment 3 suggests that children are able to represent causation when it is unintentional. However, children's tendency to use causal language more often for intended than unintended events might mean that they are more likely to infer causation when it is intended than unintended.

until the moment of a collision. We know that the force does not exist prior to the impact because the properties of the force depend on the properties of the object that is hit. For example, a car that hits a balloon will impart less force on that entity than a car that hits a bowling ball. Forces are quantities that are created at the moment of interaction. This property about forces overlaps with findings concerning the perception of causation. Since the early work of Michotte (1963), it has been noted that physical contact can have an impact on people's impressions of causation. The main overall finding is that the impression of causation is weakened in the absence of physical contact (Scholl & Tremoulet, 2000; White, 2014; Hubbard, 2013). The significance of physical contact on the impression of causation has been observed not only in adults (e.g., White, 2011, Yela, 1952), but also in very young infants (Cohen & Amsel, 1998; Leslie & Keeble, 1987; Newman et al, 2008). Related phenomena include the findings that putting a "tool" in a spatial gap makes the event seem more causal (Hubbard & Favretto, 2003; Young & Falmier, 2008) and the finding that the impression of causation is weakened as the spatial overlap between two objects grows (Scholl & Nakayama, 2002; Scholl & Nakayama, 2004; Rolfs, Dambacher, & Cavanagh, 2013). This last finding implies that the impression of causation is exquisitely sensitive to the edges of the interacting objects, not just to their overall physical distance. From a force dynamic perspective, physical contact is not necessary for causation because not all forces depend on physical contact (e.g., social forces), but physical contact is a common way in which forces can be created, and hence is predicted to serve as a valuable cue to causation (Wolff, 2008; Wolff & Shepard, 2013).

A third and final way in which forces can be created is through force redirection (Wolff, Jeon, Klettke, & Li, 2010). Force redirection occurs in the use of simple machines, such as levers, pulleys, inclined plans, wedges, screws, and wheels and axles (Correrell & Kamminga, 1990). The notion of force redirection offers an account of why certain instruments such as knives can be viewed as causers, as exemplified in the sentence in (3b). A single force vector has but one direction and magnitude. Any change to its direction or magnitude constitutes a new force. To see how forces might be created from a simple machine, consider the case of knife. A knife is a wedge. As such, it operates by converting a force applied in the direction of one edge into forces that are perpendicular to the applied force, as depicted in Figure 8. Thus, when someone cuts a cake or loaf of bread, the knife, in effect, creates two new forces perpendicular to the direction of the force from the agent.

*********** Insert Figure 8 about here ************

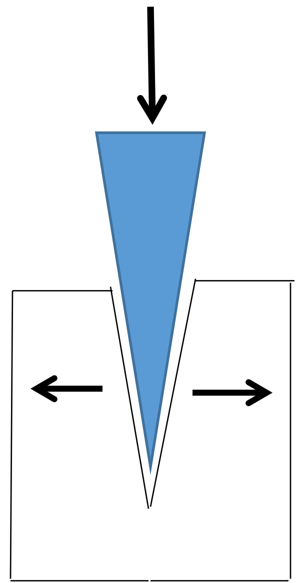

Figure 8. A downward force acting as a wedge (e.g., a knife) is redirected to create two new forces that are perpendicular to the original force.

When determining which factors in a situation are causes, people may focus on the elements of the scene that are creating new forces. Force creation through re-direction may explain why the sentences in (5) are acceptable.

5 a. The key opened the door.

b. The knife cut the bread.

c. The axe split the log.

d. The diamond scratched the glass

Crucially, not all instruments make acceptable causers in English, as shown in (6). The reason why may be because they are not construed as creating a force through re-direction (Wolff, Jeon, Klettke, & Li, 2010).

6 a. The snow shovel moved the snow.

b. The fork lifted the potato.

c. The spatula flipped the pancake

d. The broom cleaned the room.

In Wolff, Jeon, and Li (2009), we provided an initial test of this hypothesis by having participants rate sentences like those listed in (5) and (6) with respect to the affector's ability to generate its own energy. As predicted, causal sentences with high energy creation affectors were rated as more acceptable than causal sentences with low energy creation affectors. The results provide further evidence for the view that people choose the cause of a situation by identifying the causal factors that can not only be construed as causes (as opposed to enabling conditions), but also that can be construed as energy or force creators.

Mechanism. A number of studies have shown that in attempting to identify the cause of an event, people try to find the means or mechanism by which a candidate cause is able to have its effect (Ahn & Bailenson, 1996; Ahn & Kalish, 2000; Ahn, Kalish, Medin, & Gelman, 1995; Johnson & Ahn, this volume). The question of whether people attend to mechanism is important because it raises problems for many major theories of causation. According to dependency theories, a cause is a factor that makes a difference to the effect. In single events, dependency theories are expressed in the form of counterfactuals: A causes B if and only if both A and B occur, and if A had not occurred, B would not have occurred (Sloman & Lagnado, 2014). According to process theories, of which force dynamics is an example, A causes B only if there is a mechanism by which A can have an influence on B.

Walsh and Sloman (2011) provided a recent further test of the importance of mechanism. In their experiments, participants read descriptions of events in which in some cases there was an "interrupted mechanism" – for example, Frank kicks a ball, Sam moves out of its path, and the ball smashes a window. Dependency theories treat interrupted mechanisms as causes, since Sam's moving out of the path of the ball makes a difference to whether the window breaks. Process theories, on the other hand, do not treat interrupted mechanisms as causes because Sam did not transmit any force to the ball. Consistent with the predictions of process theories, participants were more likely to rate generative factors as causes (e.g., Frank kicking the ball) than they were to rate interrupted mechanisms as causes (e.g., Sam stepping out of the ball's path). Further evidence for the importance of mechanism is the many studies showing that temporal cues, and in particular temporal contiguity, have a major impact on people's judgments of causation (Greville and Buehner, 2010; Lagnado and Sloman, 2006; Lagnado, Waldmann, Hagmayer, & Sloman, 2007; McCormack, Frosch, Patrick, & Lagnado, 2014; Rottman, Kominsky, & Keil 2014; Shanks, Pearson, & Dickson, 1989; White & Milne, 1997).

For process theories, the importance of mechanism can be both motivated by and explained in terms of forces (Wolff, 2007). With respect to motivation, a force dynamic perspective encourages a local level of granularity on the analysis of causal relationships. The local nature of causal connections implies that when there is a causal connection between noncontiguous events, a reasoner assumes (in the case of physical causation) that there must be a causal chain of intermediate links to explain how forces might be transmitted or removed to bring about an effect. While a reasoner may make this assumption, in practice they usually will not know exactly how the progression of forces actually occurs. As argued by Keil and colleagues (Rozenblit & Keil 2002), people often feel as if they understand how everyday objects operate, but when they are asked to specify these operations, it becomes clear that they have little knowledge of the underlying mechanisms. Keil and his colleagues refer to this phenomenon as the “illusion of explanatory depth”.

In terms of specification, the force theory in particular provides an account of how causal chains might be formed to create a link joining a candidate cause to a particular effect. From a force dynamic perspective, the process of representing a mechanism involves the establishment of spatial connections between objects that allow for the transmission and removal of forces. A force dynamic approach to mechanism differs from a transmission view of mechanism. For example, according to Kistler’s (2006) transmission theory of causation, “Two events c and e are related as cause and effect if and only if there is at least one conserved quantity P, subject to a conservation law and exemplified in c and e, a determinate amount of which is transferred between c and e.” Kistler’s (2006) proposal builds on Dowe’s (2000) Conserved Quantity Theory. Force theories are highly related to transmission theories of causation, but they are not the same. Most notably, transmission theories are restricted to relationships in which there is a transfer of conserved quantities. But as noted by Dowe (2001), such a restriction means that transmission theories are unable to represent the notion of PREVENT and causation by omission (Woodward, 2007). Force theories, on the other hand, are able to address these phenomena because 1) they do not require that an effect occur in order for a force relationship to be present and 2) they allow for interactions in which forces are not only transmitted, but also removed.

As shown in Walsh and Sloman (2011), people seem to use this mechanistic information to select a particular factor in a situation as the cause of an effect. Mechanism, then, constitutes one last part of the answer to the problem of causal selection.

Challenges and Conclusions

Force theories address many of the phenomena associated with causal cognition. They do so by decomposing the concept of causation into factors that can be tied to both properties of the physical world as well as people’s sensory experience (Wolff & Shepard, 2013). One challenge for force theories is how such an approach might be extended to more abstract domains. Early work on this problem suggests that physical and abstract causation are understood in much the same way (Wolff, 2014), but some of the crucial tests of this hypothesis have yet to be conducted. It may be that for abstract causal relations, the mind shifts to a different kind of representational format with its own kind of combinatorial logic. On the other hand, it may be that the underlying representational substrate for abstract causation is largely the same as for concrete causation, implying that the representation of abstract causation largely preserves the properties of concrete causation (Wolff & Barbey, 2015).

As mentioned earlier, people sometimes feel that they understand how objects operate despite having little knowledge of the underlying mechanism – the "illusion of explanatory depth" (Rozenblit & Keil, 2002). Future research on force dynamics may examine a similar illusion. In so-called causal illusions (Thorstad & Wolff, 2016), people can have an initial and illusory impression of force despite the absence of physical contact, as when a magician seems to cause an object to levitate from a distance. One possible explanation for such illusions is the conflict between an initial impression of force based on perceptual cues, and a second impression based on more deliberate analysis of the underlying mechanism. Future research may examine whether dissociable cognitive processes underlie these two impressions of force.

Examination of these issues will not only further our understanding of causation, but also give us further insight into the nature of mental representation in general.

References

Ahn, W., & Bailenson, J. (1996). Causal attribution as a search for underlying mechanisms: An explanation of the conjunction fallacy and the discounting principle. *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: 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.

Anjum, R. L., & Mumford, S. (2010). A powerful theory of causation. In A. Marmodoro (Ed.), *The metaphysics of powers* (pp. 143 – 159). London, UKL: Routledge.

Barbey, A. K., & Wolff, P. (2006). Causal reasoning from forces, In *Proceedings of the 28th Annual Conference of the Cognitive Science Society* (Mahwah, NJ: Lawrence Erlbaum), 2439.

Barbey, A. K., & Wolff, P. (2007). Learning causal structure from reasoning, In *Proceedings of the 29th Annual Conference of the Cognitive Science Society* (Mahwah, NJ: Erlbaum), 713–718.

Beavers, J. (2011). On affectedness. *Natural Language & Linguistics Theory*, *29*, 1-36.

Bowerman, M. (1982). Bowerman, M. (1982). Evaluating competing linguistic models with language acquisition data: Implications of developmental errors with causative verbs. *Quaderni di semantica, 3*, 5-66.

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.

Cohen, L. B., & Amsel, G. (1998). Precursors to infants' perception of the causality of a simple event. *Infant behavior and development, 21*, 713-731.

Cotterell, B., & Kamminga, J. (1990). *Mechanics of pre-industrial technology*. Cambridge, UK: Cambridge University Press.

Copley, B., & Harley, H. (2015). A force-theoretic framework for event structure. *Linguist and Philos, 38*, 103 – 158.

Copley, B., & Wolff, P. (2014). Theories of causation should inform linguistic theory and vice versa. In B. Copley & F. Martin (Eds.), *Causation in grammatical structures* (pp. 11–57). Oxford: Oxford University Press.

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

Dowe, P. (2001). A counterfactual theory of prevention and "causation" by omission. *Australasian Journal of Philosophy, 79*, 216–226.

Dowty, D. R. (1979). Word meaning and Montague grammar. Dordrecht: Reidel.

Einhorn, H. J., & Hogarth, R. M. (1986). Judging probable cause*. Psychological Bulletin, 99*, 3‑19.

Gärdenfors, P. (2014). *Geometry of meaning: semantics based on conceptual spaces*. Cambridge, MA: MIT Press.

Göksun, T., George, N. R., Hirsh-Pasek, K., & Golinkoff, R. M. (2013). Forces and motion: how young children understand causal events. *Child Development, 84*, 1285 – 1295.

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

Greville, W. J., & Buehner, M. J. (2010). Temporal predictability facilitates causal learning. *Journal of Experimental Psychology: General, 139*, 756.

Hale, K., & Keyser, S. (1993). On argument structure and the lexical expression of syntactic relations. In K. Hale & S. Keyser (Eds.), *The view from Building 20: Essays in linguistics in honor of Sylvain Bromberger* (pp. 53–109). Cambridge: MIT Press.

Hall, N. (2000).Causation and the price of transitivity. *J. Philos. 97*, 198–222.

Hall, N. (2004). Two concepts of causation. In Causation and Counterfactuals, eds J. Collins, N. Hall, & L. Paul (Cambridge, MA: MIT Press), 225–276.

Hall, N. (2006). Philosophy of causation: blind alleys exposed; promising directions highlighted. *Philos. Compass 1*, 1–9.

Harré, R. & Madden, E. H. (1975). *Causal powers: A theory of natural necessity*, Oxford: Blackwell.

Hart, H. L., & Honoré, A. (1985). *Causation in the law*. Oxford: Clarendon Press.

Hesslow, G. (1988). The problem of causal selection. In D. J. Hilton (Ed.), Contemporary science and natural explanation: Commonsense conceptions of causality (pp. 11‑32). Brighton, Sussex: Harvester Press.

Hubbard, T. L., & Ruppel, S. E. (2013). Ratings of causality and force in launching and shattering. *Visual Cognition, 21*, 987 – 1009.

Hubbard, T. L. (2013). Phenomenal causality I: Varieties and variables. *Axiomathes, 23*, 1-42.

Hubbard, T. L., & Favretto, A. (2003). Naïve impetus and Michotte's" tool effect": evidence from representational momentum. *Psychological Research, 67*, 134-152.

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

Johnson, S. G. B., & Ahn, W. (2015). Causal networks or causal islands? The representation of mechanisms and the transitivity of causal judgment. *Cognitive Science, 39*, 1468 – 1503.

Khemlani, S. S., Barbey, A. K., & Johnson-Laird, P.N. (2014). Causal reasoning with mental models. *Front. Hum. Neurosci., 8*, 849.

Kistler, Max. (2006). Causation and laws of nature. New York, NY: Routledge.

Lagnado, D. A., & Sloman, S. A. (2006). Time as a guide to cause. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 32*, 451.

Lagnado, D. A., Waldmann, M. R., Hagmayer, Y., & Sloman, S. A. (2007). Beyond covariation. *Causal learning: Psychology, philosophy, and computation*, 154-172.

Leslie, A. M., & Keeble, S. (1987). Do six-month-old infants perceive causality? *Cognition, 25*, 265-288.

Lombrozo, T. (2010). Causal-explanatory pluralism: how intentions, functions, and mechanisms influence causal ascriptions. *Cognitive Psychology, 61*, 303–332.

Marantz, A. (1997). No escape from syntax: Don't try morphological analysis in the privacy of your own Lexicon. In A. Dimitriadis, et al. (Eds.), *Proceedings of the 21st annual Penn linguistics colloquium. Penn Working Papers in Linguistics* (Vol. 4(2), pp. 201–225).

Marmodoro, A. (2007). The union of cause and effect in Aristotle: physics III 3. *Oxford Studies in Ancient Philosophy, 32*, 205–232.

Mayrhofer, R., & Waldman, M. R. (2014). Indicators of causal agency in physical interactions: The role of prior context, *Cognition, 132*, 485 – 490.

McCormack, T., Frosch, C., Patrick, F., & Lagnado, D. (2015). Temporal and statistical information in causal structure learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 41*, 395.

McGrath, S. (2005). Causation by omission: a dilemma. *Philos. Stud. 123*, 125–148

Michotte, A. E. (1963). The perception of causality. New York: Basic Books. (Original work published 1946)

Mill, J. S. (1973). System of logic (8th ed.). In J. M. Robson (Ed.), Collected works of John Stuart Mill (Vols. VII and VIII). Toronto: University of Toronto Press (Original work published in 1872).

Mumford, S., & Anjum, R. L. (2011). *Getting causes from powers*. New York: Oxford University Press.

Nelson, Deborah G. (1993). Processing integral dimensions: the whole view. *Journal of Experimental Psychology: Human Perception and Performance, 19*, 1114–1120.

Newman, G. E., Choi, H., Wynn, K., & Scholl, B. J. (2008). The origins of causal perception: Evidence from postdictive processing in infancy. *Cognitive psychology, 57*, 262-291.

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

Pustejovsky, J. (1991). The syntax of event structure. *Cognition, 41*, 47–81.

Rappaport Hovav, M., & Levin, B. (2001). An event structure account of English resultatives. Language, 77, 766–797.

Rolfs, M., Dambacher, M., & Cavanagh, P. (2013). Visual adaptation of the perception of causality. *Current Biology, 23*, 250-254.

Rottman BM, Kominsky JF, Keil FC. 2014. Children use temporal cues to learn causal directionality. *Cognitive Science, 38*, 489–513.

Rozenblit L, Keil F. 2002. The misunderstood limits of folk science: an illusion of explanatory depth. *Cognitive Science, 26*, 521–62.

Schaffer, J. (2000). Causation by disconnection. *Philos. Sci. 67*, 285–300.

Scholl, B. J., & Nakayama, K. (2004). Illusory causal crescents: Misperceived spatial relations due to perceived causality. *Perception, 33*, 455-469.

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

Schulz, L., Kushnir, T., & Gopnik, A. (2007). Learning from doing: Intervention and causal inference. In A. Gopnik & L. Schulz (Eds.), *Causal learning: Psychology, philosophy and computation* (pp. 67–85). Oxford: Oxford University Press.

Sloman, S. A., Barbey, A., & Hotalling, J. (2009). A causal model theory of the meaning of “cause,” “enable,” and “prevent.” *Cognitive Science, 33*, 21–50.

Sloman, S., & Lagnado, D. (2014). Causality in thought. *Annual Review of Psychology, 66*, 3.1 – 3.25.

Shanks, D. R., Pearson, S. M., & Dickinson, A. (1989). Temporal contiguity and the judgement of causality by human subjects. *The Quarterly Journal of Experimental Psychology, 41*, 139-159.

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

Thorstad, R., & Wolff, P. (2016). What causal illusions might tell us about the identification of causes. *Proceedings of the 38th Annual Conference of the Cognitive Science Society*. Philadelphia, USA.

Van Valin, R. D, Jr. (1990). Semantic parameters of split intransitivity. *Language, 66*, 221–260.

Waldman, M. R., & Mayrhofer, R. (2016). Hybrid causal representations. Psychology of learning and motivation, 65, 1 – 43.

Walsh, C.R., & Sloman, S.A. (2011). The meaning of cause and prevent: the role of causal mechanism. *Mind and Language, 26*, 21–52.

Warglien, M., Gärdenfors, P. and Westera, M. (2012): Event structure, conceptual spaces and the semantics of verbs, *Theoretical Linguistics, 38*, 159-193.

White, P. A. (2006). The role of activity in visual impressions of causality, *Acta Psychologica, 123*, 166 – 185.

White, P. A. (2011). Visual impressions of force exerted by one object on another when the objects do not come into contact. *Visual Cognition, 19*, 340-366.

White, P. A. (2014). Singular clues to causality and their use in human causal judgment. *Cognitive science, 38*, 38-75.

White, P. A., & Milne, A. (1997). Phenomenal causality: Impressions of pulling in the visual perception of objects in motion. *The American journal of psychology, 110*, 573.

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

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

Wolff, P. (2007). Representing causation. *Journal of Experimental Psychology: General, 136*, 82-111.

Wolff, P. (2008). Dynamics and the perception of causal events. In T. Shipley & J. Z In T. Shipley & J. Zacks (Eds.), Understanding events: How humans see, represent, and act on events. Oxford University Press.

Wolff, P. (2012). Representing verbs with force vectors. *Theoretical Linguistics, 38*, 237 – 248.

Wolff, P. (2014). Causal pluralism and force dynamics. In. B. Copley, F. Martin, & N. Duffield (Eds.), *Forces in Grammatical Structures. Forces in grammatical structures*: *Causation between linguistics and philosophy*.

Wolff, P., & Barbey, A. K. (2015). Causal reasoning with forces. *Frontiers in Human Neuroscience, 9*.

Wolff, P., Barbey, A. K., & Hausknecht, M. (2010). For want of a nail: How absences cause events. *Journal of Experimental Psychology: General*, 139, 191 – 221.

Wolff, P., Jeon, G., & Yu, L. (2009). Causers in English, Korean, and Chinese and the individuation of events. *Language and Cognition, 2*, 165–194.

Wolff, P., Jeon, G., Klettke, B., & Yu, L. (2010). Force creation and possible causers across languages. In B. Malt, & P. Wolff (Eds.), Words and the world: How words capture human experience. Oxford: Oxford University Press.

Wolff, P., Klettke, B., Ventura, T., & Song, G. (2005). Expressing causation in English and other languages. In W. Ahn, R. L. Goldstone, B. C. Love, A. B. Markman, & P. Wolff (Eds.), Categorization inside and outside the laboratory: Essays in honor of Douglas L. Medin (pp. 29–48). Washington, DC: American Psychological Association.

Wolff, P., & Malt, B. (2010). The language-thought interface: An introduction. In B. Malt, & P. Wolff (Eds.), Words and the world: How words capture human experience (pp. 93-110). Oxford: Oxford University Press.

Wolff, P., Ritter, S., & Holmes, K. (2014). Causation, forces, and the sencse of touch. In *Proceedings of the 34th Annual Conference of the Cognitive Science Society* (Austin, TX), 1784-1789.

Wolff, P., & Shepard, J. (2013). Causation, touch, and the perception of force. *Psychology of Learning and Motivation, 58*, 167 – 202.

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: Erlbaum.

Woodward, J. (2006). Sensitive and insensitive causation. *Philosophical Review, 115*, 1-50.

Woodward, J. (2007). Interventionist theories of causation in psychological perspective. In A. Gopnik & L. Schulz (Eds.), *Causal learning: Psychology, philosophy, and computation* (pp. 19–36). Oxford, England: Oxford University Press.

Yela, M. (1952). Phenomenal causation at a distance. *Quarterly Journal of Experimental Psychology, 4*, 139-154.

Young, M. E., & Falmier, O. (2008). Launching at a distance: The effect of spatial markers. *The Quarterly Journal of Experimental Psychology, 61*, 1356-1370.

Young, H. D., & Freedman, R. A. (1999). *University physics*, 10th ed. Reading, MA: Addison-Wesley.

Zwarts, Joost. (2010). Forceful prepositions. In Vyvyan Evans and Paul Chilton (Eds.), Language, Cognition and Space: The State of the Art and New Directions. Equinox Publishing, 193–214.