Agency Intuitions in Physical Interactions

Ralf Mayrhofer (rmayrho@uni-goettingen.de)
Michael R. Waldmann (michael.waldmann@bio.uni-goettingen.de)
Institute of Psychology, University of Göttingen,
Gosslerstraße 14, 37073 Göttingen, Germany

Abstract

The question how agent and patient roles are assigned to causal participants has largely been neglected in the psychological literature on force dynamics. Based on the linguistic theory of Dowty (1991), we propose that agency is a prototype concept. We adapted Dowty's theory to account for scenarios showing physical interactions. In the standard Michotte launching scenario the ball entering the scene is usually assigned the agent role, whereas the ball that is being launched is viewed as the patient. We showed in two experiments that agency intuitions were moderated by manipulations of the context prior to the launching event. Altering features such as relative movement, sequence of visibility, and self-propelled motion tended to increase agency attributions to the patient relative to the standard scenario. We suspect that shifts in figure-ground perceptions, and intuitions about characteristics of interventions may be the overarching reason for the efficacy of the tested criteria.

Keywords: force dynamics; causal reasoning; agency; Michotte task; physical causality

Introduction

Currently there is a debate between two competing frameworks modeling causal reasoning. One prominent class is *dependency theories*, including covariation theories, counterfactual theories, and causal Bayes nets. The ontology expressed by these theories contains causal variables that either encode the presence or absence of events, facts, and properties, or different values of continuous quantities. These variables are interconnected by causal arrows that represent hidden mechanisms, and whose strength can be numerically expressed by causal strength parameters (see Waldmann & Hagmayer, in press, for an overview).

A completely different view answers the question why an observed lawfulness holds by focusing on the participants involved in a causal relation, for example Ball A and Ball B in Michotte's task, or Aspirin and a person with headache in a medical scenario. One variant of this view, *dispositional* theories of causation, would say, for example, that the ingestion of Aspirin relieves headaches because Aspirin has an intrinsic property, a disposition (or capacity or power), to relieve headaches in suitable organisms (see, for example, Gnassounou & Kistler, 2007; Mumford & Anjum, 2011).

In psychology force dynamics, an example of a dispositional account, has become increasingly popular in recent years. Pinker (2007) has argued that force dynamics is a major competitor of Bayes net theories because it allows us to model intuitions about the generative processes underlying observed covariations. One attractive feature of dispositional theories and force dynamics in particular is that these theories are capable of expressing abstract intuitions about mechanisms without requiring elaborate knowledge.

Force dynamics has been initially been developed in linguistics in the context of verb semantics (see Riemer, 2010; Talmy, 1988) but uses concepts that can be traced back to Aristotle (see Gnassounou & Kistler, 2007). Aristotle explained efficient causation as a consequence of the interaction of two entities, an agent and a patient. An agent is, according to Aristotle, a substance operating on another substance, the patient, which is suffering the process of change. The acting agent who affects the patient therefore has the *disposition, capacity* or *power* to act; and the patient has the disposition, capacity or power to undergo the agent's action.

In linguistic theories of verb semantics and argument structure verbs place constraints on the possible participants mentioned in the noun phrases. For example, in "Peter pushes Mary", "push" has two arguments, one describing an agent (Peter), the other the patient (Mary). Typically, agents are assigned the syntactic subject position. Other participant roles (also called thematic or theta roles) have been postulated but there is no agreement in linguistics about the proper list (see Riemer, 2010, for an overview). Another important semantic theory for a theory of causation is Talmy's (1988) theory of force dynamics. He argues that intuitions about the interaction of forces are an important component of our general semantic intuitions.

Using a force dynamics framework, White (e.g., 2006, 2009) demonstrated the difference between intuitive causal representations and physics by studying Michotte type launching events. In Michotte's (1963) famous demonstrations of phenomenal causality, subjects observed moving objects. For example, in a launching scenario Object X, a ball, moves towards Object Y, another ball, and touches it. This stops Object X and sets Object Y into motion at the same or a slightly lesser speed. Observers typically describe this scenario as a case in which the movement of Object Y is caused by Object X (i.e., launching). Although according to Newtonian physics the force on body Y exerted by body X is equal in magnitude but opposite in direction to that on body X exerted by body Y, observers often see Object X as the cause and Object Y as the effect (causal asymmetry). Nobody would describe the scenario as a case of Object Y stopping Object X, although this would be a legitimate description.

The impression of causal asymmetry is also reflected in judgments of force. White (2009) presented participants with different launching events, and asked them to provide estimates of the relevant underlying forces. The results showed that in such events more force is attributed to Object X than to Object Y, and that Object X is viewed as active and exerting a force on Object Y, whereas the initially stationary Object Y is viewed as inactive, exerting resistance to being moved. Thus, causal interactions are perceived as the result of the opposition between forces of agents (e.g., Object X) and resistance of patients (e.g., Object Y).

A related theory that aims at elucidating our understanding of abstract causal concepts, such as "cause", "prevent", and "enable", is Wolff's (2007) theory of force dynamics (see also Wolff & Song, 2003; Wolff et al., 2010). As in the theories of White (2009) and Talmy (1988), two entities are distinguished, which Wolff calls affectors and patients (i.e., the entity acted upon by the affector)(Talmy labels them antagonist and agonist). Force theory states that people evaluate configurations of forces attached to affectors and patients, which may vary in direction and degree, with respect to an endstate, that is, the possible result. Forces can be physical, psychological (e.g., intentions) or social (e.g., peer pressure). Causal relations are analyzed in terms of three components, (a) the tendency of a patient for an endstate, (b) the presence or absence of concordance between affector and patient, and (c) the degree to which the endstate is reached. For example, force theory would represent the singular causal fact "Winds caused the boat to heel" in terms of a patient (the boat) that had no tendency to heel (Tendency = No), the affector (the wind) acted against the patient (Concordance = No), and the result (heeling) occurred (Endstate approached = Yes).

Empirical support for the model was provided in a series of experiments in which participants made judgments about 3-D animations of realistically rendered objects (e.g., moving boats on a lake) with trajectories that were wholly determined by the force vectors entered into a physics simulator (see also Beller et al., 2009; Wolff et al., 2010; for further developments).

The Empirical Basis of Agency Intuitions

In psychological research on force dynamics the main focus has been on how causal intuitions can be predicted on the basis of configurations of forces attached to agents and patients. The assignment of the roles of agent and patient to the causal participants has typically been treated as self-evident. In Wolff's (2007) example "Winds caused the boat to heel" there is no question that the winds should be assigned the agent role because obviously they play the role of actively overcoming the passive tendency of the boat. However, one deficit of current psychological versions of this theory is that no systematic set of empirical criteria has been laid out that unambiguously motivates the assignments of the agent and patient roles.

White's (2006, 2009) theory represents progress in this regard because he has pointed out one important criterion in the Michotte tasks, movement. For example, in a typical study White presented situations in which one ball stands still, and the other moves toward this ball launching it to the other side. In this scenario the moving ball is clearly assigned the role of the agent. However, when both balls were moving the assignment was less clear, and additional assumptions had to be made for agency assignments (see White, 2012, for an extended variant of White's, 2009, theory to predict pushing vs. pulling intuitions in cases in which both balls leave the scene attached to each other).

Another example of the ambiguity of agency assignments comes from a recent study by Mayrhofer and Waldmann (2013). In this study, the alien mind reader task was used that had been introduced in the literature by Steyvers et al. (2003; see also Mayrhofer, Hagmayer, & Waldmann, 2010, for details). This scenario describes a set of aliens some of which have the capacity of picking up the thoughts of each other. In our studies the causal dependency relations were kept constant. In general a causal transmission process was described in which the thoughts of one alien (i.e., the cause) were transmitted into the heads of three other aliens (i.e., the effects). What we manipulated were subjects' assumption about the underlying dispositions responsible for the observed causal transmission. In one condition, the cause alien was assigned the role of the agent. Here the instructions stated that the cause alien has the capacity to send out his thoughts, and plant them into the heads of the effect aliens. In the contrasting condition, the effect aliens were described as the agents, being capable of reading the thoughts of the cause alien. We empirically ascertained that subjects shared our intuitions about the different assignments of the agent role. Interestingly, the results were clear-cut in the sender but ambiguous in the reader scenario. In the sender condition it was clear that the cause alien was the agent. However, in the reader scenario the intuitions did not uniquely attribute agency to the reader side, but equally divided agentive responsibility to the two sides, cause and effects. This may be analogous to the intuition that, although radios play an important part in picking up radio waves, the sending station also plays an active role.

What these results show is that it cannot always be a priori determined how the agency role is assigned, and sometimes the complementary participants may be both viewed as equally active. Hence our current goal is to empirically investigate empirical indicators of agency.

We are not the only ones who noticed that occasionally it is difficult to determine who should be assigned the agent role. For example, in the sentence "John hits Mary" John is clearly assigned the agent role. This example might suggest that grammatical subjects encode agent roles. However, in "John admires Mary", both participants play an active role so that a clear assignment is often impossible (similar with other psychological verbs, such as "mind reading")(see Dowty, 1991). Other ambiguous verbs include "buying" and "selling." They both require two active participants, and it is hard to uniquely assign the agent role.

Our experiments represent an initial attempt in a physical domain (launching events) to study factors moderating agency assignments. As a heuristic for criteria we will use Dowty's (1991) linguistic theory of the distinction between (proto-)agents and (proto-)patients. According to Dowty, agency is a prototype concept that can be assigned on the basis of a number of empirical criteria. None of these criteria is necessary (hence prototype) but the confidence of the assignment should increase the more criteria are present. In Dowty's theory the *agent* features include (a) volitional involvement in the event or state, (b) sentience (and/or perception), (c) causing an event or change of state in another participant, (d) movement (relative to the position of the other participant), (e) exists (independently of the event named by the verb). The complementary *patient* features include (a) undergoes change of state, (b) incremental theme, (c) causally affected by another participant, (d) stationary relative to movement of another participant, (e) does not exist (independently of the event named by the verb). According to Dowty, when two participants are involved in a scenario, the relative number of properties from these lists decides about the assignment of roles. If there is an impasse, multiple assignments are possible. Dowty's criteria are developed to capture semantic implications of verbs. In scene perception other cues might additionally be used, including covariation.

How can this theory be applied to launching events? Obviously some of the criteria (e.g., sentience) do not apply. The remaining relevant criteria for agency include volitional involvement, causation, and relative movement. We believe that a unifying principle behind these three criteria is provided by the intervention concept popular in both the dependency account (Woodward, 2003) and force dynamic theories (White, 2006). In fact, White (2006) believes that force dynamic intuitions are grounded in sensomotoric experiences of actors sensing resistance from the objects they are attempting to manipulate.

The prototypic agent is a human confronted with a stationary scenario that either is constant or changes in a predictable way (i.e., the ground)(relative movement). The agent's act, which is considered independent of the target system, creates a change in a variable, which in turn affects other variables (i.e., the figure)(volition; causation). Following developmental evidence (Muentener & Carey, 2010) we believe that people also make agency attributions when some of the features of the prototype are missing. For example, when no human agent is visible, the object behaving like it was manipulated by a hidden agent plays the role of the agent (e.g., one ball in the Michotte task). Apparently uncaused covariation against an invariant background has features of an intervention, which explains why the entity involved in the covariation will play the role of the agent. White's (2007) finding that the moving participant is assigned the agent role is also captured because a typical explanation of an apparently unmotivated movement is that there may be an invisible force causing it.

This theory, although derived from a force dynamic framework, is reminiscent of theories proposed within the dependency paradigm. The distinction between figure and ground is analogous to Cheng and Novick's (1991) criteria of the distinction between causes and enabling conditions. According to their theory, causal events that remain invariant within a focal set are assigned the enabler role, whereas the event covarying (i.e., changing) with the effect within the focal set is the cause. This theory does not distinguish between agents and patients, however. Another closely related theory is the intervention account of causal Bayes net theory (Woodward, 2003). According to this theory a change of a variable by a free agent qualifies as an intervention. Thus, volitional involvement, movement, and causation are hallmarks of this concept. Note that this theory is not restricted to human agents. Every change of a variable that deterministically sets the variable and has characteristics of statistical independence with respect to the target system can play the formal role of an intervention.

Although there are analogies between force dynamic and dependency theories, there are also differences. The dependency theories mainly focus on covariation and event causation, whereas other criteria we will study, including relative or spontaneous movement, are neglected.

Our main empirical strategy will be to present scenarios involving an event with two participants (e.g., two balls) but manipulate across conditions properties of the participants possibly relevant for the assignment of participant roles (agent, patient). Our aim is to show that intuitions based on the proto-agency theory predict whether a participant in a fixed scenario is assigned the agent or patient role (or both).

Experiment 1

In our experiments we employed variants of the Michotte task, which has been used as a classic demonstration for the usefulness of force theories. White (2006, 2009) has extensively studied this task, and has found a causal asymmetry effect: Agents are typically assigned greater force than patients. Another observation consistent with causal asymmetry is that the agent, for example a Ball X that is moving toward a stationary Ball Y, is typically described as causing the movement, but Ball Y is never described as stopping Ball X. Our goal is to manipulate the Michotte task in a way that either Ball X (the pushing ball) or Ball Y (the pushed ball) are more or less viewed as agents. In line with the proto-agency theory we predict that subjects differentiate between a stationary scenario (the ground) and an event that shares properties with hypothetical interventions. In Michotte's task a stationary scenario either consists of a set of balls at rest, or balls that are constantly moving in a predictable way. Given that no volitional agent (i.e., a human) is visible, other properties of causal agency apply.

As baseline condition (Condition A), we used the standard Michotte launching setup that was also used by White (2006, 2007, 2009): Ball Y is at rest in the middle of the display, Ball X is constantly moving and rolls from the left edge toward Ball Y. After contact, Ball Y moves and Ball X is at rest. In this condition Ball X should clearly be seen as the agent and Ball Y as the patient.

In three further conditions, we manipulated agency indicators for Ball Y while holding the properties of the physical interaction (i.e., the collision event) constant. Thus, in all conditions Ball Y is at rest in the middle of the screen immediately prior to the collision. When Ball X hits Ball Y, then Ball X stops and Ball Y moves (with exactly the same speed and direction as Ball X had prior to the collision) towards and then beyond the edge of the screen.

What we manipulated were the conditions prior to the invariant launching event. In Conditions B and C, we manipulated relative movement. In both conditions Ball Y enters the screen from the bottom, moves towards its position in the middle of the screen, and stops there just before it is hit by Ball X. Thus, we added movement as an agency indicator expecting an increase of agency intuitions regarding Ball Y. We suspected that observing the movement of Ball X toward the middle position might lead some subjects to infer an intention to stop Ball X. Given that this possible intention is not successful (Ball Y will be launched by Ball X), this may not fully overcome the assignment of agency for Ball X, but a difference to the standard condition may be expected.

Furthermore, we manipulated which of the balls was seen first by restricting the section of the scene that is visible to the subject. In Condition B, we hid the left hand side margin of the scene; whereas in Condition C the lower margin was hidden. Thus, in Condition B the movement of Ball Y was seen first (i.e., Ball Y is the ground); in Condition C Ball X was already seen moving when Ball Y enters the scene (i.e., Ball Y is the figure). Our goal motivating the sequence of visibility was to test whether this subtle figure-ground manipulation affects agency attributions despite the fact that the underlying physical events are identical across the two conditions. We predicted that viewing Ball Y first as the figure would increase agency attributions regarding this ball.

In Condition D, Ball Y is at rest outside the trajectory of Ball X in the lower part of the screen. Suddenly, Ball Y starts to move so that it ends up in the same position as in the other conditions. Here the constant movement of X should be viewed as stationary (i.e., ground) with Ball Y behaving like an animate volitional agent. It is well known that spontaneous movement is seen as an indicator of animacy. In this condition the intuition that the self-propelled movement of Ball Y is a result of volition should be strongest, which should lead to the strongest agency intuitions within the set of conditions.

To sum up: From Condition A to Condition D we added more and more agency indicators for Ball Y (relative movement; relative visibility; volition). According to the proto-agency theory, we expect an increasing willingness of participants to judge Ball Y as the agent in the scenarios. Of course, given that Ball Y is always eventually launched by Ball X we did not expect a complete reversal of agency assignments.

Method

Participants 39 students (27 women; mean age 23.4 years) from the University of Göttingen, Germany, participated in this experiment as part of a series of various unrelated computer-based experiments in our computer lab. Participants received either course credit, or were paid €8 per hour.

Material For each condition, we constructed a flash movie of size 760 x 760 pixels that played for 3,000 milliseconds; the first and last 400 milliseconds presented a black screen resulting in an effective movie length of 2,200 milliseconds. Ball X and Ball Y were 120 pixels in diameter; one colored in red, the other in blue. In the standard condition (Condition A), Ball Y rests in the middle of the screen such that the left most point of Y coincides with the center of the scene. After 20 milliseconds Ball X enters the scene from the left side on a horizontal trajectory with constant speed until it reaches the center of the screen (and, therefore, Ball Y) after 1,100 milliseconds. Then Ball X stops moving, and at the same time (no time lag) Ball Y starts moving with the same speed as Ball X towards the right hand side of the screen. Ball Y leaves the screen after 2,180 milliseconds. (Thus, the movie is symmetric in time and space.)

Keeping the movement shown in Condition A constant, we slightly altered the events prior to the launching event in the other conditions. We only manipulated the 800 milliseconds prior to the collision after which the movement pattern of the balls were identical across all conditions. In Conditions B and C, Ball Y entered the scene at the bottom after 20 milliseconds and moved vertically upwards until it reached its final position after 800 milliseconds. In Condition B, we covered 240 pixels of the scene's left hand side with a white panel; in Condition C, 240 pixels of the bottom were covered. Thus, in Condition B Ball X entered the scene after 700 milliseconds (whereas Ball Y was visibly moving the whole time); in Condition C Ball Y entered the scene after 840 milliseconds (whereas Ball X was visibly moving the whole time).

In Condition D, Ball Y was at rest in the lower half of the display (200 pixels above the bottom) and started moving upwards after 900 milliseconds (at the same speed Ball Y moves in Conditions B and C), and stopped at its final position after 800 milliseconds (i.e., movement time of 300 milliseconds). This sudden, apparently self-propelled movement was expected to suggest a volitional intervention into the trajectory of Ball X.

For counterbalancing purposes we additionally generated seven more movies per condition by rotating the scene by 90°, 180°, and 270°, respectively, and switching colors of the balls yielding 4 x 2 = 8 movies per condition (i.e., in sum 32 movies).

Procedure We presented each subject with all 32 movies in random order. After seeing a movie (self-paced), we requested participants to select one of four sentences (presented in randomized order) as the best description of the scene:

1. The red ball launched the blue ball.
2. The blue ball stopped the red ball.
3. The blue ball launched the red ball.
4. The red ball stopped the blue ball.

Note that only two of the sentences actually described what was seen in the movie. If a subject selected one of the two nonsensical sentences we coded the answer as an error.

Design and Prediction We recoded subjects' responses according to the color coding as "X launched Y" vs. "Y stopped X" (plus error), and aggregated the eight color/rotation versions to align with a consistent X/Y assignment. We expected an increasing selection rate for "Y stopped X" and a decreasing selection rate for "X launched Y", respectively, from Condition A to Condition D.[1]

Results and Discussion

Fig. 1 shows the average selection rates for the two relevant scene descriptions across the four agency conditions. In line with previous research, Condition A revealed a strong preference in selecting Ball X as the agent (94.9% vs. 3.9%). As predicted, selecting Ball X decreased from Condition A to Condition D, $F_{3,114}$=24.0, p<.001, η^2=.39. The preference for seeing Ball Y as the agent increased analogously, $F_{3,114}$=22.9, p<.001, η^2=.38. The average error rate was 2.6% and did not significantly differ across conditions, $F_{3,114}$<1.

Experiment 1 clearly demonstrates that agency intuitions are grounded in empirical indicators of agency, and confirmed the proposed proto-agency theory. However, it could be argued that the forced-choice format forced people to choose one description even when their intuition was in line with the symmetry assumptions of Newtonian mechanics. This argument does not explain why on average the choices did not even out, but we still wanted to replicate the results of Experiment 1 using a more unrestricted response format.

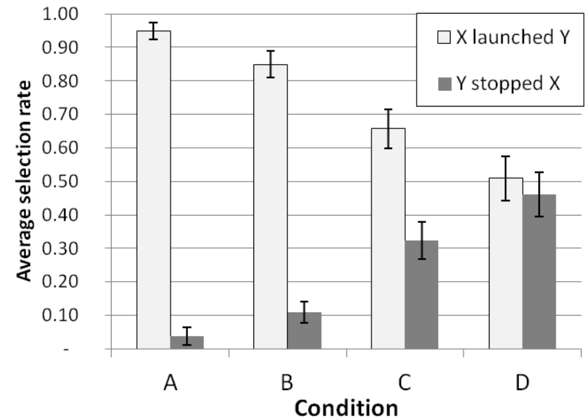

Figure 1: Results of Experiment 1. Error bars indicate standard error of the means.

Experiment 2

The goal of Experiment 2 was to replicate the findings of Experiment 1 with an unrestricted response format that allows subjects to express that they see both alternative sentences as valid descriptions of the scene. To accomplish this goal we presented subjects in Experiment 2 with rating scales that allowed them to judge the appropriateness of the scene descriptions independently.

Method

Participants A new set of 34 students (23 women; mean age 23.4 years) from the University of Göttingen, Germany, participated in this experiment using the same design as in Experiment 1.

Material and Procedure We used the same set of 32 movies and the same procedure as in Experiment 1. Instead of a forced choice decision between scene descriptions, we presented subjects with the two sentences (adapted to the respective color version), and requested them to rate how well the sentences describe the scene using two separate rating scales ranging from 0 ("not appropriate at all") to 10 ("highly appropriate"). Both sentences and rating scales were presented on a single screen; the order of the sentences was counterbalanced within subjects.

Design and Predictions We aggregated subject-wise across color/rotation conditions, which yielded a 4 (agency condition) x 2 (Ball X vs. Ball Y) within-subjects design with agency ratings as dependent measure. Since we expected decreasing ratings for Ball X and increasing ratings for Ball Y, we predicted an interaction between agency condition and the rated ball (X vs. Y).

Results and Discussion

Fig. 2 shows the results for Experiment 2. As expected, the ratings for Ball X were higher as for Ball Y in Condition A with a decreasing trend for Ball X and an increasing trend for Ball Y from Condition A to Condition D. This pattern led to a significant interaction, $F_{3,99}$=23.7, p<.001, η^2=.42. Across conditions, Ball X received higher agency ratings than Ball Y, $F_{1,33}$=34.3, p<.001, η^2=.51, reflecting the fact

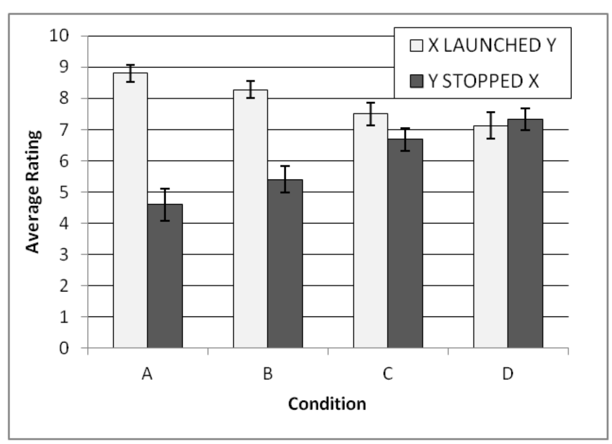

Figure 2: Results of Experiment 2. Error bars indicate standard error of the means.

[1] Note that both measures are not independent of each other; selection and error rates add up to 1.

that the salient end of the scene (Ball Y's leaving) overall dominates agency intuitions.

The results of Experiment 2 resemble the results of Experiment 1 closely. Although the difference between rating Ball X as the agent vs. rating Ball Y as the agent is much smaller in Condition A compared to Experiment 1, the overall pattern (decreasing ratings for Ball X and increasing ratings for Ball Y from Condition A to Condition D) was replicated, and showed that the findings were nor restricted to specific response formats.

General Discussion

In contrast to dependency theories, force and dispositional theories of causal reasoning incorporate the distinction between agents and patients in causal interactions. The principal focus of research motivated by dispositional accounts was on how force configurations predict causal judgments, whereas the assignment of the agent and patient roles has largely been treated as self-evident. Various studies in both linguistics and psychology have shown, however, that role assignments are not always clear-cut. Occasionally it may even be necessary to assign the agent role to multiple causal participants.

Based on the linguistic theory of Dowty (1991), we proposed that agency is a prototype concept with multiple criteria, none of which necessary for the role assignment. We adapted this theory to account for physical interactions (e.g., Michotte type launching events). In the standard Michotte launching scenario the ball entering the scene (and launching the other ball) is typically assigned the agent role, whereas the ball that is being launched is viewed as the patient. We showed that agency intuitions are moderated by manipulations of the context prior to the launching event. Altering scene features, such as relative movement, sequence of visibility, and self-propelled motion tended to increase agency attributions to the patient relative to the standard scenario.

A unifying principle underlying these criteria may be that they all tend to lift the patient into the foreground (i.e., into the figure role), and appear to suggest some kind of volitional intervention. Intervention seems to be a central concept unifying dependency and force theories, although the criteria for determining agency are different in these two frameworks. We realize that our experiments just represent a first step. Future studies will have to go beyond launching scenarios to arrive at a more complete theory of agent/patient assignments.

Acknowledgments

We wish to thank Sven Ritter for help in preparing the experiments. This research was supported by a research grant of the Deutsche Forschungsgemeinschaft (WA 621/22-1) that is part of the priority program "New frameworks of rationality" (SPP 1516).

References

Beller, S., Bender, A., & Song, J. (2009). Weighing up physical causes: Effects of culture, linguistic cues and content. *Journal of Cognition and Culture*, *9*, 347-365.

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

Dowty, D. (1991). Thematic proto roles and argument selection, *Language*, *67*, 547-619.

Gnassounou, B. & Kistler, M.(2007). Introduction. In M. Kistler & B. Gnassounou (Eds.)(2007). *Dispositions and causal powers*. Aldershot: Ashgate.

Mayrhofer, R., Hagmayer, Y., & Waldmann, M. R. (2010). Agents and causes: A Bayesian error attribution model of causal reasoning. In *Proceedings of the Thirty-Second Annual Conference of the Cognitive Science Society* (pp. 925-930).

Mayrhofer, R., & Waldmann, M. R. (2013). *Agents and causes: Dispositional intuitions as a guide to causal structure*. Manuscript submitted for publication.

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

Muentener, P., & Carey, S. (2010). Infants' causal representations of state change events. *Cognitive Psychology*, *61*, 63-86.

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

Pinker, S. (2007). *The stuff of thought: Language as a window into human nature*. New York: Viking.

Riemer, N. (2010). *Introducing semantics*. Cambridge: Cambridge University Press.

Steyvers, M., Tenenbaum, J. B., Wagenmakers, E-J., & Blum, B. (2003). Inferring causal networks from observations and interventions. *Cognitive Science*, *27*, 453-489.

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

Waldmann, M. R., & Hagmayer, Y. (in press). Causal reasoning. In D. Reisberg (Ed.), *Oxford Handbook of Cognitive Psychology*. New York: Oxford University Press.

White, P. A. (2006). The causal asymmetry. *Psychological Review*, *113*, 132-147.

White, P. A. (2009). Perception of forces exerted by objects in collision events**.** *Psychological Review*, *116*, 580-601.

White, P. A. (2012). Visual impressions of pushing and pulling: The object perceived as causal is not always the one that moves first. *Perception*, *41*, 1193-1217.

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

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

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.

Woodward, J. (2003). *Making things happen: A theory of causal explanation*. Oxford: Oxford University Press.