Journal of Experimental Psychology:
Learning, Memory, and Cognition
1999, Vol. 25, No. 1, 116–136

0278-7393/99/$3.00

Inferences Through Imagined Actions: Knowing by Simulated Doing

Daniel L. Schwartz and Tamara Black
Vanderbilt University

People make simple physical inferences by acting on objects. They might, for example, tilt a container to determine its chances of spilling. Five experiments examined whether people can also draw physical inferences by taking simulated actions. The basic experimental task involved 2 glasses of different widths but equal heights. People imagined that the glasses were filled to the same level with water and answered whether they would spill at the same or different angles. When asked explicitly, people were usually wrong; but, when they closed their eyes and tilted each glass until the imagined water reached the rim, they correctly tilted a narrow glass farther than a wide one. These experiments dissociated simulated actions from both propositional inferences and visual imagery. The results suggest a new emphasis on the role of motor activity in drawing inferences and address issues related to joining naive intuitions and explicit understandings.

As people operate in their physical environments, they constantly make inferences about simple physical events. These events may be proximal and immediately under their control, as in the case of deciding how tightly to hold a glass so it will not slip. Also, the events may be more distal, as in the case of inferring how far a glass can tilt before its water will spill. In an effort to understand how people make the more distal form of inference, scientists have explored either people's naive descriptions of events or their ability to see or visually imagine events. One might call the former "knowing by description" and the latter "knowing by viewing." These are both powerful methods of inference. However, neither reflects the often interactive and bodily quality of people's everyday physical inferences. People, for example, might tilt the glass to help make their inference. In this case, one may say that people are "knowing by doing."

When knowing by doing, people draw inferences about the world through their actions. The question addressed by this article is whether people can also draw inferences through a mental simulation of knowing by doing. Can they, for example, mentally simulate tilting a glass and in this way determine how far it can tilt before spilling? To explore this question, we first examine a few intuitive cases of knowing by simulated doing to see how they compare and contrast with knowing by doing, knowing by description, and knowing by viewing. Our goal is to develop a handful of properties that can establish simulated doing as a distinct form of inference. Afterward, we return to the context of tilting a water glass to test whether these properties apply. In the General Discussion, we develop the implications of our discoveries for theories of imagery and for topics in the traditions of naive physics and embodied cognition.

Daniel L. Schwartz and Tamara Black, Department of Psychology and Human Development, Peabody College, Vanderbilt University.

John Rieser, John Bransford, Mary Hegarty, and Gedeon Deak have been thoughtful and constructive sources of ideas. The suggestions and questions of Nancy Cooke and Timothy Hubbard were extremely helpful. Parts of Experiment 2 were reported in April 1995 at the annual meeting of the American Educational Research Association in New Orleans, Louisiana.

Correspondence concerning this article should be addressed to Daniel L. Schwartz, Box 512-GPC, Vanderbilt University, Nashville, Tennessee, 37203. Electronic mail may be sent to Dan.Schwartz@Vanderbilt.edu.

Knowing by Simulated Doing Versus Knowing by Doing: Representation

The study of knowing by doing has typically occurred within the tradition of ecological psychology (e.g., Gibson, 1962). Turvey, Solomon, and Burton (1989), for example, have demonstrated that without looking at a rod, people can determine its length by wielding it. Moving the rod generates torque and inertial information that specifies its length. The authors claimed that wielding generates sufficient perceptual information to solve the problem without representational enrichment. In this case, inferences are the directly perceivable results of action.

One property that distinguishes knowing by doing from a simulation is the use of representation. One would not normally consider the wielding of the rod to be a simulation. For it to be a simulation, people need to represent aspects of the situation that are not perceptually available.

People can represent the same physical situation in different ways and can use different problem-solving strategies. We want to identify the representational properties that seem particularly central to the strategy of simulated doing. Consider, for example, a chain of five horizontally linked gears. Assume that the gear on the far left rotates clockwise, what will the gear on the far right do? One way to solve this problem is to apply a parity rule which states that all odd gears turn one way and all even gears turn the other. Another way to solve the problem is to model the gears, perhaps using one's hands to represent the position and motion of each interacting gear (Schwartz & Black, 1996b). It is the gestural representation and not the rule representation that satisfies most people's intuitions for what counts as a simulation (e.g., Clement, 1994).

Knowing by Simulated Doing Versus Knowing by Description: (Imagined) Actions in Real Time

By far, the most studied form of physical inference is knowing by description. A description is an assertion or proposition about the world; for example, "all the odd gears turn the same way." It is a belief about which one could say "true" or "false." Most scientific discourse is descriptive. Science education research and computational models of qualitative physics typically examine how people learn or know the world through descriptions. By examining how more accurate and refined descriptions replace old ones, it is possible to map the evolution from naive to sophisticated science understanding and to make educationally relevant prescriptions (e.g., White, 1993). A short list of representations that reside under the umbrella of descriptions includes qualitative predicates (e.g., de Kleer & Brown, 1984; White & Frederiksen, 1990), presuppositions (e.g., Vosniadou & Brewer, 1993), naive theories and beliefs (e.g., Clement, 1983; McCloskey, Caramazza & Green, 1980), analogies (e.g., Gentner et al., 1997), and various forms of local and global rules (e.g., Hegarty, Just & Morrison, 1988; Williams, Hollan & Stevens, 1983).

Despite the importance of knowing by description, problem solving that depends exclusively on description does not reflect a key aspect of "doing." This is because descriptions involve neither actions nor imagined actions. Rules about gears, for example, do not involve simulated actions, whereas gestural models do. When knowing by doing, people use actions to draw an inference. For example, they may turn one gear to see how a neighboring gear responds. Likewise, simulated doing should depend on actions, even if those actions are only imagined.

Of course, imagined actions cannot generate feedback from distal objects. Nevertheless, they may still tap into unique sources of knowledge. Rizzolatti et al. (1988), for example, have reported premotor neurons that discharge to a visual stimulus but only when that stimulus is presented in the context of a possible action (even if no action is ever taken). Conceivably, such action-oriented neurons can be recruited to the process of drawing a physical inference.

More symbolically, there are computer models that use simulated actions to draw inferences (Gardin & Meltzer, 1989; Schwartz & Black, 1996a). A relevant property of these models is that their representations are timing responsive; they respond to real-time rates of change. Suppose, for example, that a model represents a tilting glass that contains either molasses or water. As the model simulates turning the glass, it generates a succession of timing "beats" that indicate the rate the glass changes. The molasses and the water representations respond to these beats. Perhaps the molasses requires two beats to take one step of change toward the rim of the glass, whereas the water requires only one. In any event, without the beats, the water and molasses representations do not manifest effects of viscosity, and they cannot help make inferences.

Humans may also have timing-responsive representations. Mental maps of self-position, for example, appear to be timing responsive. Rieser (1990) asked people to walk on a treadmill being pulled across a field by a tractor. When the tractor pulled faster than the treadmill turned, people saw themselves moving through the field faster per step than usual. Over a period of minutes, people automatically recalibrated the coupling between their mental maps and their stepping. When taken off the treadmill and asked to walk to a point blindfolded, they did not walk far enough. Their mental map of their position in the environment changed too fast in response to the beat of each step, and people thought they were covering more ground than they actually were.

Timing-responsive representations can help explain why people spontaneously simulate actions when they do not have adequate descriptive knowledge (Clement, 1994; Schwartz & Hegarty, 1996). People, for example, use gestural simulations prior to inducing descriptive rules or when their rules fail (Schwartz & Black, 1996b). Compared to descriptions, actions generate the inference-relevant beats upon which timing-responsive representations depend. For example, when people model a slow gear rotation with their hands, they presumably generate a slow series of beats. In contrast, descriptions have an arbitrary relation to real time; a description of a slow action can itself be slow or fast. So, by this scheme, the beats of (imagined) actions can naturally harness timing-responsive knowledge representations in a way that descriptions cannot.

It is not straightforward to make a compelling empirical demonstration that imagined actions are responsible for an inference. It is necessary to show that descriptive representations are not tacitly responsible for the inference. This is problematic at a theoretical level because descriptive representations can be used to characterize most anything (Anderson, 1978). When solving the gear-chain problem, for example, people's hand gestures may simply be outward manifestations of underlying descriptions (cf. Pylyshyn, 1973). Perhaps people first tacitly describe what their hands should do, and then they let the hands perform according to description.

Given the ever-present theoretical possibility that a description is behind a simulation, it is important to show empirically that people do not have descriptions that could lead to simulation outcomes (e.g., Decety & Michel, 1989). Krist, Fieberg, and Wilkening (1993), for example, showed that when 5-year-old children push an object off a ledge to hit a target, they correctly reduce their pushing speed to compensate for increases in ledge height; yet, when asked, children believe that the higher an object's release point, the faster it must be pushed to reach a target. This study captures the type of evidence that indicates that knowing by doing is different from knowing by description. The current experiments tried to gather this type of evidence for simulated doing.

Knowing by Simulated Doing Versus Knowing by Simulated Viewing

Another salient feature of simulated doing is its perceptual quality. In the gear case, people's gestures create a series of spatial displacements that resemble what people would see if watching the gears. In contrast, when people use a parity rule, they reason about numerical properties (e.g., odd and even). If they do evoke a perceptual trace it is secondary to the problem solving. The similarity between models in the imagination and actual visual input is frequently invoked as a criterion for demonstrating imagery (Finke, 1985; Kosslyn, 1976; Perky, 1910). There is a large body of evidence indicating that people can use visual imagery to reason about shape, position, change of perspective, and navigation. For example, people can mentally rotate an object until they "see" whether it matches the appearance of a target object (Shepard & Cooper, 1986). If one follows the similarity between imagery and perception, one might call visual imagery "knowing by simulated viewing." People simulate what they would see.

Metric Information

Characterizing simulation as a form of visual imagery captures an important quality of knowing by doing. Actions, just like visual perceptions and their representational analogs, move through intermediate states. They are continuous in the sense of being sensitive to metric information. A solution requires less transformation (and time) if there are shorter distances and fewer intermediate states. Therefore, evidence of simulated doing should indicate that people's representations are sensitive to metric properties of the referents. For example, people should not only be able to infer the qualitative outcomes of putting two gears side by side (directions of motion), they should also be able to infer the effects of using different sizes of gears (velocities of motion).

Causal Propagation

Visual imagery captures some aspects of simulated doing. It excludes others. Problems about physical interactions, for example, often concern the forces that cause changes. Whether or not forces are perceptible to vision (Gilden, 1991; Hecht, 1996; Hubbard, 1995b), they are sensed directly through the action of the haptic system (Turvey, 1996). Therefore, it seems like a mistake to confine questions about simulations to the modality of vision and its representational analogs. Moreover, one would expect simulations to model the causal propagation of forces to draw inferences, as in the case of modeling one gear turning the next, and then the next, and so on. Visual imagery does not always have this constraint. For example, one may imagine rotating a figure clockwise or counterclockwise to see if it matches a target figure. In contrast, because of the directional nature of forces, simulated doing may not have the freedom to go forward or backward and still draw inferences about physical events (cf. Freyd, 1983). Hegarty (1992), for example, found that people's "mental animations" are superior when they can simulate a pulley system forward from the precipitating cause compared to when they must move backward from the resulting effect (i.e., the pull on a rope vs. the resulting rise of a block).

Although causal propagation may be a necessary feature of simulated doing, it is not an ideal feature for distinguishing simulations from visual imagery. This is because visual imagery can mimic the forward propagation of causes. To distinguish simulated doing from visual imagery, there should be an additional property.

Inferences in Action Not Vision

Another way that imagery research may not capture the essence of simulated doing is that it often emphasizes the visually apparent consequences of actions rather than the actions themselves (e.g., Finke, 1980). When people model the gears, they take actions by letting their hands act out the behaviors of the gears. Accompanying this activity, there is a visual perception or imagination that, to some extent, resembles gears. In tasks that explicitly require people to match an image against a visual percept or a visual memory, the image's appearance is the key to completing the task (e.g., Rock, Wheeler, & Tudor, 1989; Shepard & Cooper, 1986). But, for imagined physical events, visual appearance may be derivative of other key inferences. For example, the gear gestures, in and of themselves, do not indicate how people know to turn their hands in opposite directions; they only show it. Moreover, through knowing by doing, people can express their knowledge without making a judgment based on visual appearance. People, for example, might express their inferences through an action or virtual action.

Relative to visual imagery, appearance may play a diminished role in simulated doing. This suggests a possible dissociation between the two on the basis of people's awareness of visual appearance. Assume that simulated doing derives in part from motor representations. There is evidence that motor imagery does not always code actions from the perspective of vision. It can also code actions from what Jeannerod (1994) referred to as the "first-person" perspective and which includes the feeling of executing an action. Moreover, motor representations can operate outside of consciousness (Bridgeman, 1992). For example, if people's reaching movements toward a visual target are perturbed, they adjust so rapidly (less than 100 ms) that they cannot possibly be using an awareness of the perturbation to guide their behaviors (e.g., Pelisson, Prablanc, Goodale, & Jeannerod, 1986). The motor system can operate independently of an awareness of visual appearance.

We do not know of motor research on awareness that involves imagining a distal physical event, as would be the case when tilting a glass and imagining the movement of its water. Perhaps for physical inferences like this, people would monitor the goal state in a visual image. They might, for example, imagine the water "rising" to the rim of a tilting glass. However, as they tilt their hands, they may not visually imagine or be aware of the angle of the glass that objectively controls this rising. The real-time beats that arise from the tilting action may be enough to drive the imagined water transformation, regardless of the glass's visual appearance. By this account, visual imagery can be a component of simulated doing, but there are also components that manifest

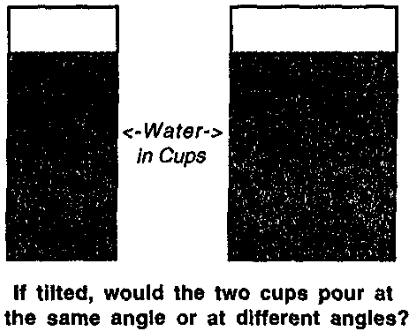

Figure 1. Example stimuli for the basic inferential task.

themselves through actions, of whose consequences people can be visually unaware.

Experimental Operationalization

We have proposed five properties that may help characterize inferences made by simulated doing: They involve *representations* of perceptually unavailable information. They depend on (imagined) *actions in real time.* They represent *metric information.* They model the *causal propagation* of effects. They can generate inferences in action and *without complete visual awareness.* To see if these ideas extend to a different context from the gear example, we tested them in the context of people reasoning about the behavior of water in a glass.

To set the stage, consider the following problem. There are two glasses of identical height, and they have lines indicating equal levels of imaginary water. The only difference between the two glasses is that they have different diameters. Figure 1 provides a schematic of two glasses that fit this description. If the glasses had actual water, would they start pouring at the same or different angles? If the glasses would start pouring at different angles, which glass would pour first? When presented to 8 pilot participants, only 1 individual described the correct outcome. A second group of 8 people held one of the glasses. They closed their eyes and tilted the glass until they thought the imagined water would just reach the rim. After the experimenter recorded the angle of tilt, the individuals repeated the process with the other glass. In this condition, all 8 people were implicitly correct. They correctly tilted the narrow glass farther than the wide glass. These simulation results are impressive, in part, because gestural problem solving has primarily been documented when people's hands stand for the object or concept of interest (e.g., the gear gestures; Alibali & Goldin-Meadow, 1993). In the current case, however, people's gestures helped them to represent water even though their hands were not used as surrogates for that water.

People may have solved the problem by simulating the effects of actions over imagined physical constraints. These constraints may have included the knowledge that liquids conform to their containers, that objects are "pulled" in the direction of gravity, and that volume is conserved. As the glass tilted, the conformity and gravity constraints caused the imagined water to "spread" against the side of the glass closest to the ground. The conservation of volume constraint ensured that when the water spread toward the rim on one side of the glass, it retreated along the other side by an equivalent amount. As a result, the midpoint of the water surface became an axis separating the advancing and retreating water. The radius between the axis and the glass side is shorter in the thin glass than in the wide one. Figure 2 shows that for the same degree of tilt, the edge of the water advances less for a small radius than a large one. (As an analogy, imagine rotating both a pencil and a yardstick 90°; the tip of the yardstick traverses a greater linear distance than the tip of the pencil.) As a result, people tilted the thin glass farther because it took longer for the water to advance in response to their rate of tilting. In the following experiments, we used this task to demonstrate various combinations of the five hypothesized properties of knowing by simulated doing described above.

Experiment 1

The first experiment attempted to establish that people can make the correct tilting inference through imagined action even though they do not know the correct answer descriptively. For generality, the experiment consisted of three substudies, each using a different glass shape and a different sample of people. The rectangle study used a rectangular glass that could be tilted widthwise or lengthwise. The cone study used two identical cones, one open at the flared end and one open at the tapered end. The cylinder study used two household glasses that differed in diameter. In each case, people held a single glass in their favored hand and closed their eyes. They tilted the glass counterclockwise in the

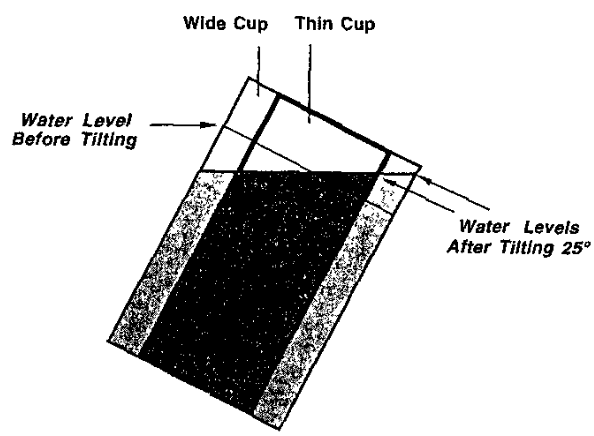

Figure 2. A cross-section of a thin and wide cup tilted simultaneously. The wide cup and the superimposed thin cup have the same water levels when upright. At the same degree of tilt, the water has swept a greater distance along the side of the wide cup than along the side of the thin cup. (See the text for further explanation.)

picture plane until the imagined water just reached the rim (clockwise if left handed). The process was repeated with the second glass. Participants received no feedback and never saw their tilts. After the two tilts, they saw the glasses and made an explicit judgment as to whether the two glasses (or rectangle orientations) would start pouring at the same or different angles. The answers are that a wide cylinder pours before a thin cylinder, that a flared cone (wide opening) pours before a tapered cone, and that a tilt of the long axis of a rectangle pours before a tilt of the short axis (see Figure 3).

Method

Participants

For the rectangle study, 25 (16 women, 9 men) faculty and graduate students volunteered. For the cone study, 20 participants (14 women, 6 men) from the larger community were paid. For the cylinder study, 16 undergraduates (7 women, 9 men) participated for partial course credit.

Materials

Rectangular cup. An 8.9-cm tall plastic cup had a rectangular cross-section of 7.2 cm × 5.5 cm. To indicate the water level, a thin strip of tape circumscribed the cup 3.3 cm below the rim. When filled to the tape, water touched the rim of the cup at 45° with a tilt of the long axis and 55° for a tilt of the short axis.

Cone cups. Each plastic cone had a 2.3-cm radius for the tapered end and a 7.0-cm radius for the flared end. The cones each had a height of 8 cm and a slant of 30°. The flared cone was open on the wide end, whereas the tapered cone was open on the thin end. Each cone had a line indicating an imaginary water level 3.0 cm below the opening. The water reached the rim of the tapered cone at a 45° tilt of the central axis and the flared cone at a 25° tilt.

Cylindrical cups. The thin and wide glasses had 1.6-cm and 4.2-cm radii, respectively. Both were 10 cm tall, with the water level indicated 2.5 cm below their rims. Water poured at 60° in the thin glass and 40° in the wide glass. The thin glass was an Irish coffee shot glass with a heavy base and a small handle near the bottom. The wide glass was a crystal coffee cup with a large handle and a slight curve at the base.

Design and Procedure

With vision, participants grasped a cup with their thumb on the water line. They were told to close their eyes and tilt the cup in the picture plane until the imaginary water just reached the rim in their imaginations. After they indicated they had finished, an incline protractor measured the tilt of the cup before they let go. The flat edge of the protractor was placed against the cup; the protractor's internal plumb indicated the angle of the cup with respect to gravity. The cup was then taken from the participants, they opened their eyes, and the process was repeated for the second cup. Participants received no feedback about their tilts. Tilting order was counter-balanced across participants. After the tilting phase, participants saw both cups simultaneously. We asked participants, "Do you think the water pours out at same or different angles for each cup [rectangle orientation]?" counterbalancing the order of the words *same* and *different*. If people answered "different," they said which would have to be tilted farther.

Results

Nearly all participants correctly tilted the cup with the narrower cross-section of rim farther than they tilted the wider cup. Figure 3 shows the average tilts, and Table 1 shows the percentage of individuals who tilted the narrow cross-section farther. Despite this accurate tilting, Table 1 shows that people's verbal judgments were not very accurate. Participants were accurate above chance only for the cones; $Z = 1.97$, given three choice alternatives. Throughout, a significance level of $p < .05$ has been adopted.

People's explicit beliefs, right or wrong, were unrelated to their tilts. A between-subjects variable was created by separating those who made a correct explicit judgment from

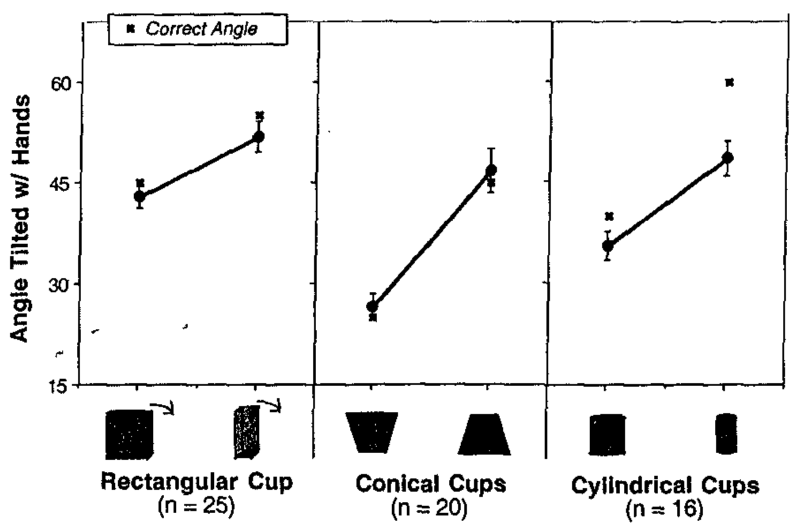

Figure 3. Average tilts for three glass shapes (Experiment 1). Error bars represent the standard error.

Table 1
Distribution of Participants' Explicit and Simulation Inferences for the Three Cup Shapes

Shape of cup	Explicit judgments (%)			Tilting behavior (%)		
	Narrow[a]	Same[b]	Wide	Narrow[a]	Same[b]	Wide
Rectangle	28.0	68.0	4.0	84.0	12.0	4.0
Cylinder	18.8	37.5	43.8	100.0	0.0	0.0
Cone	55.0	30.0	15.0	100.0	0.0	0.0

Inferences about which cup diameter must be tilted farther to pour.

[a]Narrow is the correct answer. [b]Tilts within 2° of one another are considered the same.

those who did not. Each person's pair of tilt angles served as dependent measures that captured the within-subject effect of glass diameter. For the rectangle, there was an effect of diameter on tilting, $F(1, 23) = 32.0$, $MSE = 26.89$, but no Diameter × Belief interaction, $F(1, 23) = 0.31$. For the cones, there was an effect of diameter, $F(1, 18) = 66.6$, $MSE = 60.69$, and no Diameter × Belief interaction, $F(1, 18) = 0.0$. Similarly for the cylinders, there was an effect of diameter, $F(1, 14) = 36.7$, $MSE = 22.0$, and no Diameter × Belief interaction, $F(1, 14) = 0.21$. These statistics indicate that the results of people's earlier tilts did not influence their later judgments, and they more generally reflect the independence of description and simulated doing, even though the description immediately followed the correct simulated doing.

Discussion

When tilting the glasses, people's simulations had enough resolution that 84% of the participants differentiated two objective tilts that varied by only 10° for the rectangular cup, and 100% of the participants differentiated objective tilts of 20° for the cones and cylinders. In addition to their relative accuracy, participants were reasonably accurate in absolute terms. The only cup that yielded an absolute error over 5° was the thin cylinder. Experiments 3 and 5 looked more closely at why the tilts were underestimated for the thin cylinder.

People did not use their tilting activity to inform their subsequent explicit judgments. Their explicit judgments were usually incorrect, and they were statistically unrelated to their tilting performance. People did judge the cones correctly above chance. However, their justifications were not about the movement of water or the end point of the tilting. They were about the slant of the cone either complementing or negating the angle of tilt. Thus, these people were probably not reasoning on the basis of their prior imagery of the water.

People may not have been able to capitalize on their prior tilting because their eyes were closed and haptic traces are fleeting (Milner & Goodale, 1995; Posner, 1967). Participants would need to remember one tilt with enough precision and perseverance to compare it to a subsequent tilt that varied by a small amount. Most likely, people did not try to do this. This is because they may not have been conscious of the tilt angle, at least with respect to the angle the glass makes with the horizon (cf. McAfee & Proffitt, 1991). Experiment 5 formally documented this lack of awareness.

People may have incorrect beliefs about this problem because they have never monitored a tilting cup from a perspective that would help them notice width effects explicitly. In everyday activity, the important "spill comparison" is about the level of water within a glass and not about the width between two glasses. Moreover, the tilting competency appears to be in place by the time children reach age four (Black & Schwartz, 1996); consequently, there is little need for explicit noticing anyway. Given the early competence and the further lack of an explicit catalyst to monitor the effects of glass width, it is reasonable that people never develop descriptions about the effect of glass width on water behavior.

Experiment 2

The preceding results provide the initial demonstration that simulations and descriptions are dissociable for this task. The explicit judgment task, however, does not reveal enough information about the character of people's descriptions to secure the point. Perhaps the three ordinal alternatives (same time, thin first, wide first) inhibited attention to important metric characteristics of the problem and caused the incorrect judgments (e.g., Schooler, Ohlsson, & Brooks, 1993). If true, then perhaps another task could show that people have correct descriptions that could be covertly responsible for the accurate tilting. To address this concern, Figure 4 shows a task that avoids the necessarily limited, and perhaps suggestive, range of judgment responses. On the left of the figure is a standard glass with its water level indicated. On the right is one of several target glasses, in this case narrower and shorter. The task is to draw the water level in the target glass so that both glasses need to tilt the same amount for the water to reach their rims. In the following substudies, people solved numerous problems. Each problem used the same standard glass with varied water levels. The target glasses varied in width, height, and proportional-

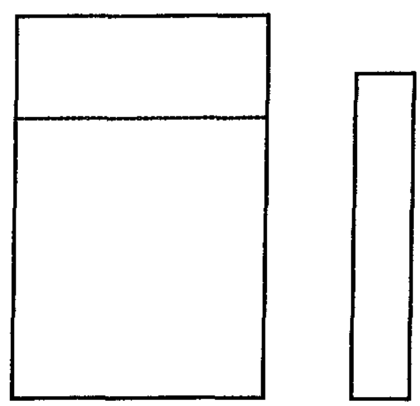

Figure 4. "Where should the water level be drawn in the thin glass so the two glasses pour at the same angle?"

ity to the standard's dimensions. In Experiment 2A, participants simply solved the problems. In Experiment 2B, individuals thought aloud as they solved a subset of the problems.

The two studies were intended to illuminate people's reasoning when they do not use simulations. Some people may try to use a simulation to solve this task, but pilot work indicated that people cannot simultaneously model the dynamics of two water-filled cups. Sixteen people tilted a board that held a wide and thin glass of imagined water that poured at 40° and 60° respectively. They stated that they could not attend to both glasses simultaneously, they tilted the board to an average of 20°, and fewer than 20% correctly indicated that the imagined water reached the rim of the wide glass first. Therefore, even if people tried to use a simulation, it was unlikely to be the basis of their solutions for the current task.

Method

Participants

Seven female undergraduates participated in Experiment 2A as partial fulfillment of course requirements. Four male and 5 female undergraduates volunteered for Experiment 2B.

Materials

For Experiment 2A, each student received a packet of 33 randomly ordered problems. The standard glass, on the left of each page, was always 11.4 cm tall × 5.1 cm wide. The standard had its water level at a high position (1.0 cm to 2.5 cm below the rim), middle position (5.7 cm), or low position (8.9 cm to 10.4 cm). On the right of the page, the target rectangle varied from 4.4 cm to 17.1 cm in height and from 1.3 cm to 7.3 cm in width. Figure 5 shows the 11 different sizes. Each target was coupled with the standard for

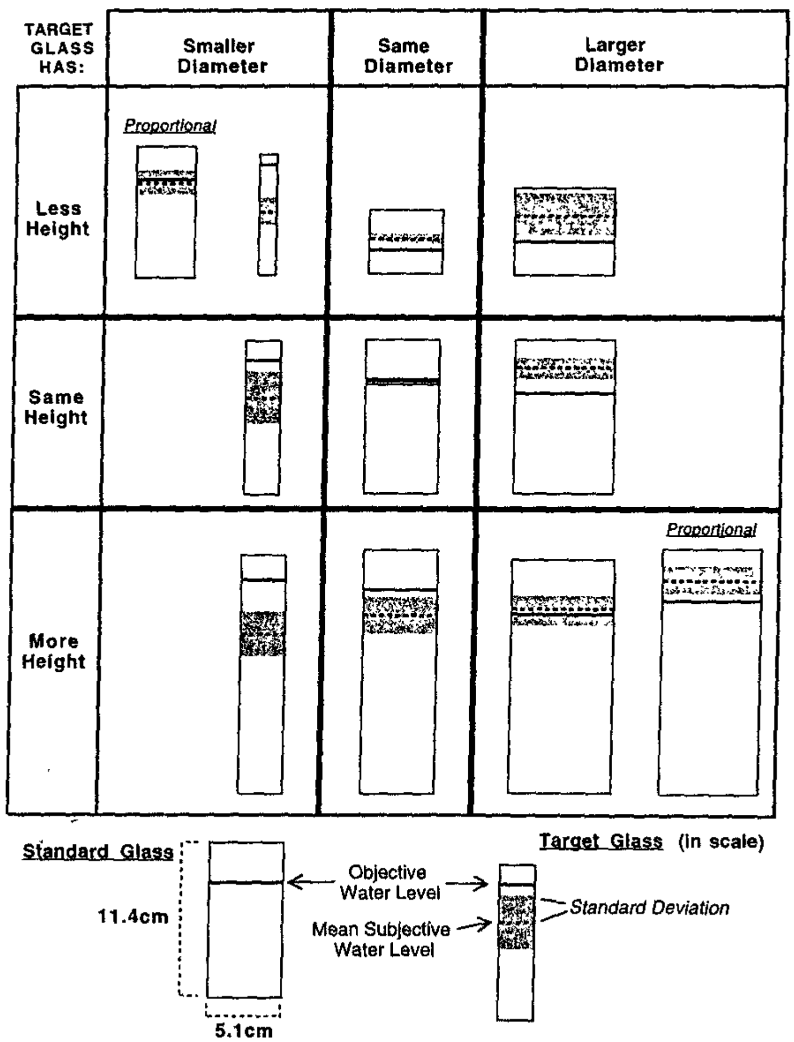

Figure 5. Answers for the draw-a-line task (Experiment 2A).

each of the three water levels. Experiment 2B used the subset of 11 problems with the high water level.

Design and Procedure

Experiment 2A may be thought of as 3 × 3 × 3 within-subject design. One variable was the width of the target glass relative to the standard (thinner, same, wider). The second variable was the height of the target glass (shorter, same, taller). The third variable was the water level in the standard glass (high, middle, low). The two target glasses that had height-to-width ratios proportional to the standard were included to illuminate features that people used to solve the problem; they were not part of the experimental design. Experiment 2B eliminated the water level variable. The protocol coding scheme is described in the *Results*. Students in Experiment 2A were tested as a group and completed the 33 problems in 10 min. Students in Experiment 2B individually completed a videotaped, think-aloud procedure that took about 30 min.

Results

Experiment 2A

Figure 5 presents the results for each target glass when collapsed across the three water levels. On each target glass, the solid line represents the average of the true water levels. The dotted line represents the average position of the participants' lines relative to the true water level. The shaded region shows the standard deviation across individuals. As may be seen, people were consistently wrong.

The discrepancy between each drawn line and the objective water level was measured and signed; placements below the true level received negative values and placements above received positive values. Each individual generated 27 measures based on the Width × Height × Water Level variables (the proportional glasses are treated below). There was a main effect due to the width of the target glass, $F(2, 4) = 114.7$. This main effect was in the opposite direction of the true qualitative relationship and applied to all 7 individuals. People drew the water level above the standard line in a wide target glass, whereas it should have been drawn below. They also drew the water level below the standard line in a thin glass, whereas it should have been above. The main effect of height was marginal, $F(2, 4) = 7.3$, but there was a Width × Height interaction, $F(2, 4) = 230.7$. This interaction reflected the tendency to displace the line upward for a wide glass and downward for a tall glass. In combination, these effects cancel for a tall and wide glass but compound for a tall and thin glass. There was no main effect of water level, $F(2, 4) = 3.5$. There was a Water Level × Width interaction, $F(2, 4) = 88.4$; all other interactions were nonsignificant. Low water levels led to greater negative errors for thin glasses and greater positive errors for wide glasses.

Figure 5 indicates that the target glasses with height-to-width ratios proportional to the standard led to different response patterns from the nonproportional glasses. The proportional glasses led people to focus on different features. An individual, for example, might attempt to make the water to no-water ratio equivalent in the target and standard glasses, because the large proportional glass made the area ratios salient. A stepwise regression determined which features and combinations of features best predicted the distance the line was drawn from the bottom of the target glass. Orthogonal vectors for width, height, and water level were constructed and multiplied by one another to create two- and three-way interaction terms. The interaction terms represent higher order relationships within the target glass. The Width × Height interaction, for example, can represent the total area of the target glass, and the Height × Water Level interaction can represent the area below the water line. Significant interactions indicate that some people compared areas and ratios among areas.

All seven vectors entered the regression model significantly and accounted for 73% of the variance in how far the line was drawn from the bottom of the glass, $F(7, 220) = 84.5$, $MSE = 6.0$. The water level was predicted to be the most influential variable, because it could serve as an anchor for the water level placements, $\beta = 0.33$, $t(220) = 9.1$. The most influential variable, however, was the width of the glass, $\beta = 0.55$, $t(220) = 15.3$. (The bivariate correlation between the drawn water level and the diameter was $r = .66$.)

Experiment 2B

The quantitative data of Experiment 2B replicated the diameter effect found in Experiment 2A. The crossed within-subject variables of height and diameter and a between-subjects variable of gender created a 3 × 3 × 2 multivariate design (the proportional glasses were excluded). The only significant effects were for the diameter of the glass, $F(2, 6) = 26.1$, Hotellings $T^2 = 8.7$, and the height of the glass, $F(2, 6) = 5.2$, Hotellings $T^2 = 1.73$. There was no Diameter × Height interaction in this experiment, $F(4, 4) = 3.5$, Hotellings $T^2 = 0.35$. This difference from Experiment 2A may be due to the extra time spent explicitly reflecting on each problem and previous answers. The overall effect of gender was negligible, $F(1, 7) = 0.0$, $MSE = 1.3$, and did not enter into any interactions (all $F\text{s} < 1$).[1]

The protocols indicate high strategy variability. They also indicate that individuals' strategies were irrelevant, because people made the same drawing error regardless of their strategy. To formally examine the different strategies people used to solve the problems, a post hoc analysis yielded nine categories that were mutually exclusive and represented a conservative portrayal of the variability in strategies. Categories, criteria, and representative quotes are shown in Table 2.

---

[1] In all of our experiments involving the water pouring task, there were no significant gender effects. This contrasts with the Piagetian water level task, which normally reveals gender effects (e.g., Liben & Goldbeck, 1984). An anonymous reviewer made the important observation that the current results may imply that gender effects are specific to declarative or solely visual tasks (e.g., the Piagetian task) and may not generalize to real-world motor tasks (e.g., pouring task). Because of our small sample sizes and because we did not administer the Piagetian task, we are not in a position to make strong claims. Future research dedicated to this issue, however, may support such claims.

Table 2
Main Strategies People Used in Experiment 2B

Label	Criterion	Sample quote
Imagine	Participants stated they were imagining, or they physically turned the page.	I am trying to imagine the two glasses at the same time.
Width	Participants noted width as a justification.	Is this narrower than the ones before, so should the line be down a little lower?
Height	Participants noted height as a justification or the distance the water had to travel along the side of the glass.	If the glass is shorter, then it's going to need more water.
Area/volume	Participants reasoned about the projected area or volume of the glass or water.	Maybe it's the volume in that pouring space, between the top of the glass and the water level.
Ratio	Participants reasoned about the relative ratios of space within or across glasses, or both.	I'll draw the line halfway because this glass looks like half the size of this glass.
Identity	Noting identity of standard and target	Identical, so put an identical level.
Physical	Participants noted force-based features of problems typically relating water mass and velocity.	(a) It's lighter, so it should move faster. (b) There's more mass and force pushing the water to the edge.
Angle	Participants explicitly noted the angle the water would make with the side of the glass.	If I look at the angle between the corner and the top of the water, that should represent the angle that cup would be at when you tilt it, because water will always be parallel with the ground.
Guess/no reason	No justification or an explicit guess.	

A primary individual coded protocols noting each strategy occurrence for each problem. A secondary coder analyzed the protocols of three randomly selected participants. The two coders had 92% agreement and exhibited no systematic differences.

Individuals used an average of 3.6 (SD = 1.0) different strategies. Even within a single problem, individuals switched between strategies an average of 1.8 times (SD = 0.6). Across subjects, each problem elicited a minimum of four different strategies, except the identical target. Figure 6 shows the frequency of strategy for each glass.

To examine the relationship between strategy use and drawing errors, we looked at the final strategy that individuals mentioned before they drew the line. We analyzed the problems that varied the width of the glass because these problems were consistently yielding the reversal of the true relationship. Figure 7 shows that regardless of strategy, people tended to draw the line closer to the rim for a wide glass than for a narrow glass (a reversal of the true relationship). To assess the effect of strategy statistically, we treated an individual's final strategy as an independent variable crossed with glass width. The position of the line was the dependent measure. (The ratio category was collapsed into the area category because the ratio strategy only occurred twice as a final strategy and used area ratios in each case.) The data were not nested under subject to avoid many empty cells. The results show the usual main effect of glass diameter for where people drew the line, $F(1, 58) = 18.9$, $MSE = 1.1$, but they do not show an effect of strategy, $F(6, 58) = 1.2$, or Diameter × Strategy, $F(6, 58) = 0.86$.

Although not significantly different, the angle strategy sometimes led to different line placements from the other strategies. Two individuals eventually discovered an accurate geometric solution that depended on drawing a diagonal from the midpoint of the original water level to the rim of the glass. The diagonal represents the tilt needed to put the water at the rim. If one reproduces this diagonal on the target glass, it is simply a matter of drawing the target's water level so that it is bisected by the diagonal. Compared to the other strategies, the angle strategy had a more articulated basis and was applied relatively consistently once found. Nonetheless, this explicit strategy conflicted with whatever caused the consistent errors in line placement. One individual who had discovered and twice applied an angle-based theory switched strategies nine times on a subsequent problem with the thinnest target. The other individual "cheated" when drawing the water level on a thin glass. He worked out the angles, placed a finger where the diagonal would intersect with the water, and then, somewhat secretively, drew the water line much lower. Evidently, people have a strong bias that is difficult to overcome.

Discussion

There is a history of dispute concerning whether imagery-like results are actually the result of tacit descriptions (Finke & Freyd, 1989; Kosslyn, 1976, 1980; Pylyshyn, 1973, 1981). To investigate this issue, we conducted two experiments, looking for evidence that descriptive knowledge, explicit or implicit, was covertly responsible for the accurate tilting. With respect to explicit descriptions, Experiment 2B showed that people do not have a stable naive theory in the sense that it drives them to reason over theoretically important attributes using theoretically specific strategies (e.g., McCloskey, 1983). People rapidly switched attention

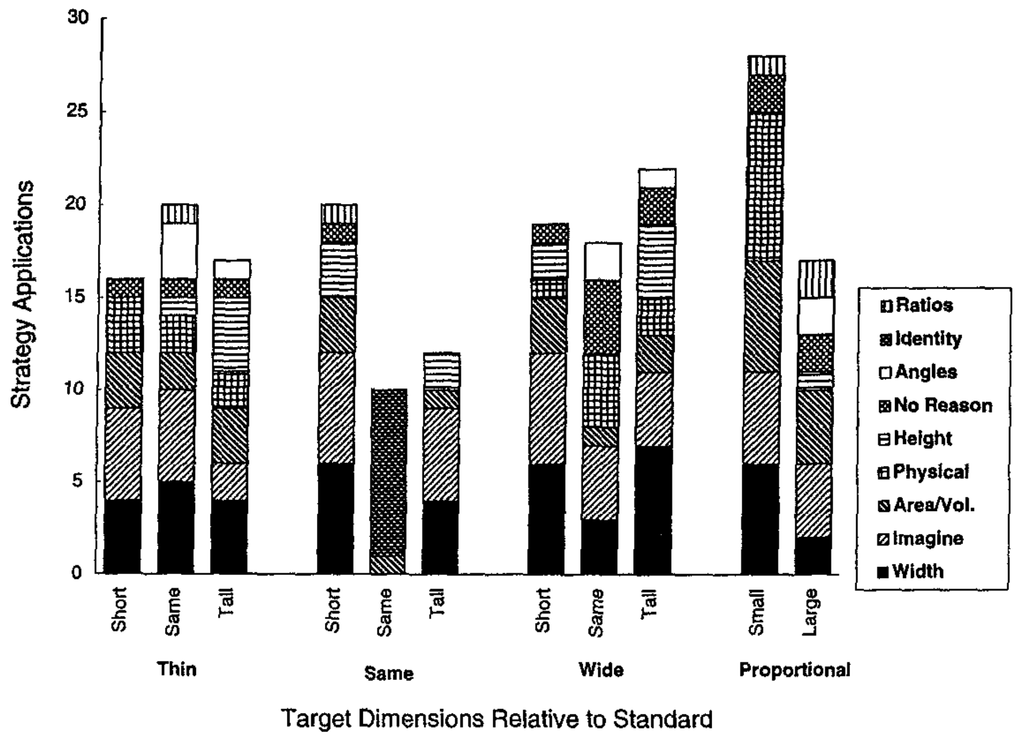

Figure 6. Strategies used for each target glass (Experiment 2B). Vol. = volume.

among problem features and strategies. It seems they were moving through various problem attributes hoping they would find a satisfactory solution (Cooke & Breedin, 1994). Unlike other demonstrations of unstable naive reasoning that have depended on isomorphic situations that have different surface features or task demands (e.g., a pendulum and projectile; Kaiser, Jonides, & Alexander, 1986), the current problems simply adjusted a few parameters. It seems unlikely that such unstable explicit beliefs could be responsible for the consistently accurate tilting.

With respect to implicit or tacit descriptions, the studies showed that people make a stable error when drawing the line that is in the opposite direction of the true qualitative relationship. In contrast to situations in which consistent errors indicate stable naive understandings (Brown & Burton, 1977), the stable errors here cannot be attributed to a naive theory; people's explicit beliefs varied and were unrelated to the errors they made. Because there is not yet a satisfactory explanation for the genesis of this systematic error, it is considered a bias. In tandem, the demonstrations of an incorrect implicit bias and of unstable explicit beliefs suggest that people's simulation performance and their descriptive reasoning stem from different knowledge formats.

Experiment 3

Given the elimination of description as the cause of the successful tilting, we turn our attention to the properties of simulated doing. One property is that simulations should be sensitive to metric information. Many computational models of physical inference have tried to limit computational complexity by primarily considering qualitative relationships (for a review, see Forbus, 1988). Qualitative relationships typically capture moments of change, as in the case of a ball switching from falling, to hitting the ground, to rebounding. Actions, however, often require metric information. If people wanted to catch the ball, they would need to know its velocity and size so they could gauge how fast to close their hand. The current experiment used four different levels of imagined water to determine whether people's simulations include metric information. If they do, then people should tilt the same glass farther when the water level is lower.

The experiment also tested whether overt motor action is a necessary feature of simulated doing. In Experiment 1, people had to represent the water but did physically turn each glass. In this experiment, participants needed to represent the glass and their actions over it.

The paradigm used in this experiment was as follows: People were first shown a glass with its indicated water level. They were then told to close their eyes and to imagine the glass turning until the water reached the rim. During this imagining, they did not make any hand motions whatsoever. Because people were not outwardly modeling their thoughts with gestures, it was necessary to develop a way to gain

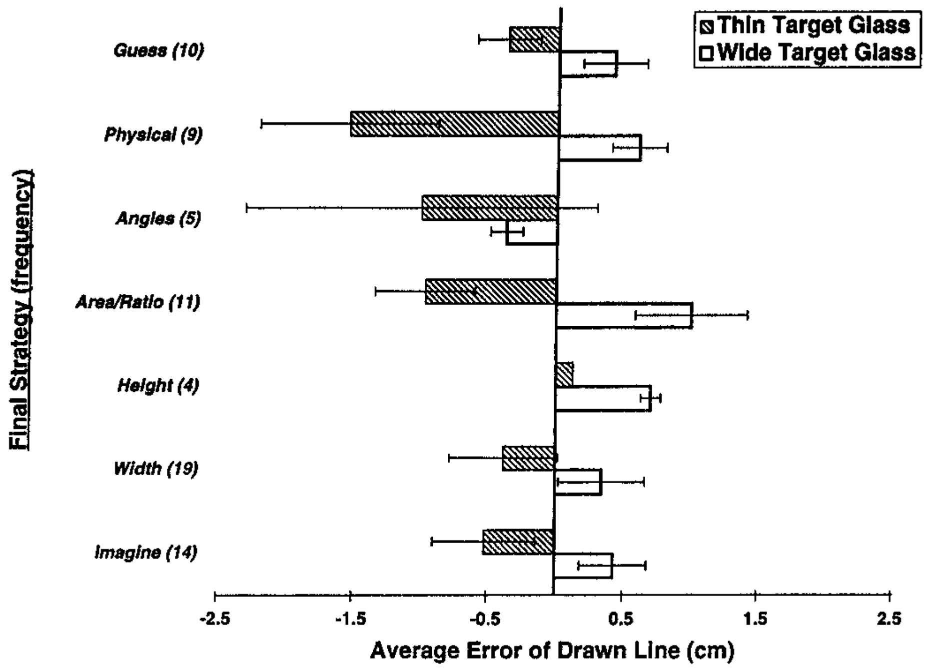

Figure 7. Drawing errors for each final strategy and glass width (Experiment 2B). Error bars represent the standard error.

access to their simulations. The approach taken here was to have people say when they were done with their mental simulation. At that point, they were asked to raise their hand to indicate the angle of the imagined glass by putting a "lid" on it. Although Experiment 1 suggested that people do not naturally attend to the angle of the glass, participants in this study were explicitly asked to attend to this information after they completed their tilting. The *Method* section describes this in more detail. For now, it is important to clarify our claim about visual awareness. Our claim is not that people cannot "view" the imagined appearance of their simulations if asked to. Rather, our claim is that people may not naturally do this. For example, they may be so focused on the imagined water reaching the lip of the glass that they may not pay attention to the angle of the glass. However, if asked, they may retrieve this information if it is contained in the image. Therefore, our speculation was that individuals would indicate a larger tilt for a thin glass, compared with its corresponding wide glass, and they would increase their tilts as the water levels decreased.

Method

Participants

Fifteen graduate and undergraduate students participated for financial compensation or partial course credit. There were 9 women and 6 men.

Materials

Eight cylindrical glasses were 15.5 cm tall. The four thin glasses were 3.1 cm in diameter. The four wide glasses were 9.5 cm in diameter. Each glass had a single water level indicated at 1.5, 3.0, 4.5, or 6.0 cm below the rim.

Design and Procedure

The crossed diameter and water level variables created eight tilt measures. During the imagined tilting phase, participants saw only one glass at a time in random order and closed their eyes during the imagined tilting. After the tilting phase, individuals made an explicit judgment about a pair of thin and wide glasses.

Individuals neither handled the glasses nor moved their hands in imitation of a tilting glass. Instead, they were shown a glass at eye level for a few seconds. They closed their eyes and in the picture plane tried to simulate the tilting glass and water solely in their imaginations. If individuals stated that they lost the image, they were allowed to start over. To report the imagined tilt, individuals positioned a flat stick as though placing a cap on top of the tilted glass. Individuals completed the imagined tilt, and, with eyes still closed, raised the stick in front of themselves and lowered it so it would lie flat on the imagined opening of the glass. The angle of tilt was measured by placing an inclined protractor on the stick. This capping procedure was used because pilot work demonstrated that tilting the stick to indicate the angle of the side of the glass interfered with the image of the glass (cf. Wohlschläger & Wohlschläger, 1998). To practice capping before the experimental task, individuals looked at a pretilted glass, closed their eyes, capped the glass, and then opened their eyes to receive feedback.

Five to 10 training trials were usually sufficient for people to reach accuracy within 5°.

Results

Simulating four pairs of glasses and explicitly reporting their tilt angles did not improve the explicit judgments compared to Experiment 1. Only 2 of the 15 participants correctly stated the wide glass would pour first (9 said same time, and 4 said thin first). In contrast, Figure 8 shows that people could complete the simulation without overt action. Participants' tilts had the correct relationships across glass width. Moreover, they were sensitive to the metric property of water level. The thin glasses were tilted farther than the wide glasses for each water level, $F(1, 14) = 54.5$, $MSE = 26.4$; and individuals made larger tilts for lower water levels, $F(3, 42) = 105.7$, $MSE = 24.2$. This latter effect was located in the linear component of the Tilt × Water Level relation, $t(42) = 13.6$, $SE = 1.7$, with a marginal curvilinear component, $t(42) = 1.9$, $SE = 1.2$. There was no Diameter × Water Level interaction, $F(3, 42) = 1.8$, $MSE = 12.6$.

Figure 8 shows that participants did not tilt the glasses as far as they should, especially the thin glasses. There are many possible explanations including the way they were asked to translate their imagery into an external representation, their tendency to imagine water in a tilted glass at an incline rather than parallel to the ground (e.g., Pascual-Leone & Morra, 1991), an overestimation of the rate of water movement, superior discrimination in the midrange of the turn (e.g., Turvey, 1996), eight trials causing a "satiation" effect, or a general conservativeness about spilling water. Experiment 5 visited this issue again.

Discussion

The current results demonstrate that the simulations represent metric information and are not purely qualitative: People's angle of imagined tilt increased as the water level lowered. Moreover, the results show that people can conduct the simulation without motor action and concurrent perceptual stimulation. There was anecdotal evidence that the lack of bodily motion made the task more difficult. Compared with Experiment 1, individuals required more time and restarts to complete the task. Presumably, there are tasks (and individuals) for which the extra information and constraint provided by motor activity are necessary to sustain the imagery. At the same time, there are tasks for which no amount of bodily action helps—for example, depicting two glasses with water simultaneously. The precise nature of the task complexity that requires or exceeds bodily action is not forthcoming from these experiments.

Even though the angle information from each tilt was made salient by the reporting method, participants still did not use this information in their explicit judgments; they made the usual judgment errors. Given the results of Experiment 2, it seems likely that people have such a strong bias that whatever imagined information is available, it cannot sway them. In fact, Schwartz and Hegarty (1996) found that 85% of their experimental participants chose their "head over their hands" when told that their explicit judgments disagreed with their hand motions. Despite the evidence that people do not have stable beliefs about this task, people still prefer those insecure beliefs to the results of their simulated doing. The basis for this preference is not clear, and the preference is not a general phenomenon. Elsewhere, we have documented that people will rely on simulated actions when they do not have explicit beliefs (Schwartz & Black, 1996b). The safe conclusion here is that the "look" of the glass in the imagination is not carrying much evidential force in people's explicit reasoning. This complements the claim that people do not always use the visual appearance of imagined actions and their consequences when knowing by simulated doing.

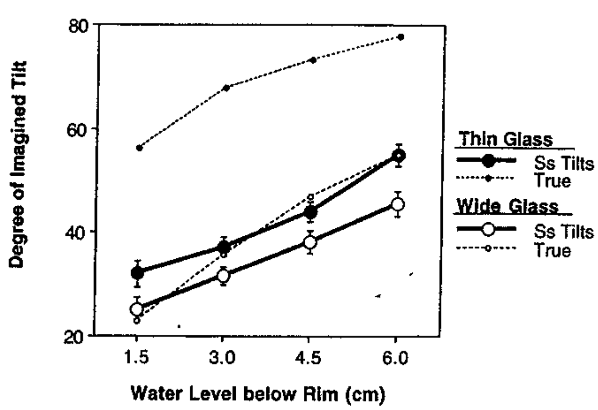

Figure 8. Tilt angles when participants did not use their hands (Experiment 3). Error bars represent the standard error. Ss = subjects (participants).

Experiment 4

The preceding experiment demonstrated that people can complete the water pouring task without concurrent motor activity and that they represent metric (continuous) properties. In this regard, simulated doing is similar to visual imagery. In the following experiments, we began to explore how they may differ. One way the pouring task is unlike many visual imagery tasks is that people need to coordinate the physical behavior of one imagined object (the glass) with the behavior of another (the water). The current experiment examined whether a representation of causality and its propagation is an important component of how people solve this problem of coordination.

One may see the value of this demonstration by considering an alternative solution based on geometric relations rather than causal ones. Imagine that people grasp the glass with their right forefinger at the water level. They tilt the glass counterclockwise so that their forefinger rotates upward while the leading edge of the rim (the edge on the opposite side of the glass where the water will eventually pour) rotates downward. They stop tilting at the point where the plane between their forefinger and the leading edge of the glass is parallel to the ground. By this strategy, people will always tilt a thin glass farther than a wide one, and they can do so without ever simulating the causal behavior of the water. Their geometric reasoning simply pantomimes causal reasoning. For example, they could use the same basic strategy if the glass started in the upside down position—simply tilt the glass rim upward until the key features are put into geometric alignment. The current experiment examined whether people are solving the water pouring task on the basis of geometric relations or whether they model causality.

In this study, people were told, as usual, that their task was to tilt the glass to the point where the imagined water reaches the rim. Before they grasped the glass, however, the glass was turned upside-down, and they were told to tilt the glass upward to indicate the position. If people use a geometric strategy like the one outlined above, they should still be able to complete the upside-down task. On the other hand, if simulations of physical events follow causal propagation, the task modification of going backwards against gravity should block their ability to conduct a simulation, and they will exhibit inaccurate tilting performance.

Method

Four male and 12 female student volunteers participated. The materials were the same cylindrical glasses from Experiment 1. In the forward tilt phase, participants tilted the thin and wide glasses downward from the upright position as usual. In the backward tilt phase, people tilted the thin and wide glasses upward from the upside-down position. The participants viewed a glass upright for a few seconds. Afterwards, the glass was turned upside-down. They were told that they should turn the glass to the position where the water would just reach the rim of the glass (if it were really being tilted downward). They were allowed to grasp the glass with any hand orientation, although 14 of 16 grasped the glass with their right thumb pointing toward the ground to facilitate a clockwise movement in the picture plane (counterclockwise if left-handed). As always, people closed their eyes when tilting the glass. The orders of the tilting phases and the presentation of the wide and thin glasses were counterbalanced.

Results

Figure 9 plots the forward and backward tilts for the two glasses. Each line represents a single individual by connecting his or her wide and thin tilts within each condition. A solid line indicates the individual correctly tilted the thin glass farther than the wide one, and a dotted line indicates the individual did not. In the forward tilt condition, everyone tilted the thin glass farther than the wide glass. In the backward condition, 10 of 16 individuals did show a greater tilt for the thin glass, but only 4 individuals tilted the thin glass more than 5° farther than the wide glass (compared to 14 of 16 in the forward condition). Glass diameter and tilt direction created a 2 × 2 within-subject design with angle from upright as the dependent measure. Indicating a different relationship between the tilts of the thin and wide glasses in each condition, the Direction × Diameter interaction was significant, $F(1, 15) = 17.3$, $MSE = 31.2$. There was a main effect for tilt direction, $F(1, 15) = 28.6$, $MSE = 214.6$, indicating that people tended to position the glass farther from the upright position when going backward as compared to forward (as usual, people underestimated in the forward direction). There was also a main effect for diameter, $F(1, 15) = 16.1$, $MSE = 51.3$, although this effect was driven by the forward tilt results. When examining the backward tilts in isolation, the thin and wide glasses did not have significantly different tilts, $F(1, 15) = 0.2$, $MSE = 62.0$.

One interpretation of the results is that participants simply had difficulty with the backward task because it was unusual. If this were the case, then one would expect the correlations between the thin and wide tilts in the backward condition to be quite low because of an increase in error variance. Table 3 provides the correlations among the four tilts. The tilts of the glasses show significant thin–wide correlations within each tilt condition but not across tilt conditions. In other words, people's errors in the backward condition were consistent and not reflective of general problems. The lack of correlation across conditions suggests that individuals engaged different processes in each condition.

Discussion

The forward-tilting condition told the usual story. In the backward-tilting condition, however, the results were normally distributed around equivalent tilts for the thin and wide glass. Starting from the upside-down position made it difficult or impossible for people to conduct their simulations. One may argue that the upside-down movement did not prevent the simulation, but instead, it entailed a longer and more unusual motion that simply led to more errors in implementation. If true, one should predict that the two trials in the backward condition would have had increased error variance and a low intercorrelation. This prediction was not supported.

The results show that action is not sufficient for solving this task. Instead, the action needs to simulate the "normal" causal propagation of events. The results also show that imagery theories that only consider geometric information cannot explain the current behaviors. These theories cannot predict that one direction of tilt should be substantially more difficult than the other, just as they do not predict that rotating an object clockwise or counterclockwise should make a difference. Participants, however, were strongly influenced by the direction they tilted the glasses.

Many participants mentioned they had difficulties with gravity. Several people said that as soon as they started to turn the glass from upside down, the water ran out in their imagination. Several other individuals mentioned that they imagined ice so the water would not pour out of the overturned glass. They said that they let the ice thaw once the glass had tilted for awhile. This suggests that the people tried to handle the upside-down condition by changing the situation so they could model a forward flow of events. Unfortunately, they could not overcome their perception of gravity. This put them in a difficult position because they needed to ignore gravity (or make the water ice) so the water

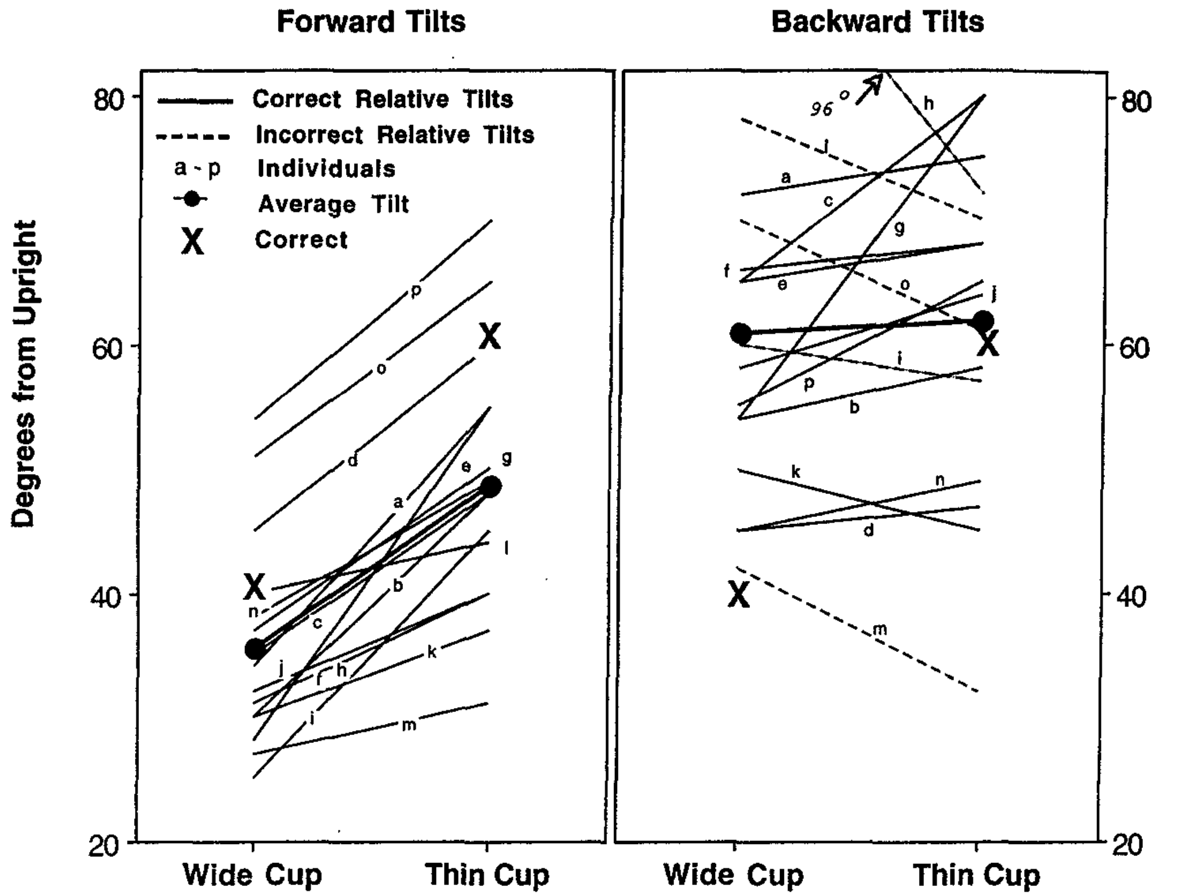

Figure 9. The effect of tilting the glasses backward from the upside-down position (Experiment 4). Each line connects an individual's wide and thin tilts within a condition. The letters track the individuals across the forward and backward conditions.

would not run out of the glass, but at the same time, they needed gravity (or liquidity) so the water would conform to the upside-down glass as it was tilted. It is an open question why nobody represented an "upwards gravity" to help the solve the problem. One possibility is that under other conditions people could represent gravity (e.g., Hubbard, 1997), but for this task there was too much interference from the perception of gravity. Another alternative is that people cannot represent an analog of gravity because there is rarely a need; gravity is always perceptually available. Lackner (1992), for example, documented that under weightless conditions, people's "attempts to imagine being in a particular orientation always failed to bring about the perception of being in that orientation" (p. 811). An interesting test of these two alternatives might be to ask people to complete the water pouring task in weightless conditions to see if they can correctly simulate the water's behavior without gravity's presence.

Table 3
Correlations Between Tilts

Tilt direction and glass width	1	2	3	4
1. Forward thin	—	—	—	—
2. Forward wide	.79	—	—	—
3. Backward thin	.25	.14	—	—
4. Backward wide	−.01	−.01	.68	—

Experiment 5

The preceding experiment differentiated the water simulation from geometric accounts of visual imagery by showing that the simulation depends on causal propagation. The current experiment examined a second possible difference between doing and viewing. In visual imagery tasks, people are presumably aware of how their solution would look if it were an actual percept instead of an image. In fact, this is exactly what people are often expected to determine—whether their image corresponds to a visual state of affairs or a perceptual memory. With simulated doing, however, people may not need to be aware of how their solutions look. Lederman, Klatzky, Chataway, and Summers (1990), for example, argued for direct haptic apprehension. They showed that by handling a given object, people can represent and make inferences about that object without recoding the haptic information into a visual image. Similarly, in the pouring task, people may not recode the angle of the wrist and glass into a visual image (unless they are requested to do so, as in Experiment 3). Even though the glass angle is the objective control variable for the movement of water in a glass, people may not attend to the angle in their simulations. Perhaps they coordinate the glass and water on the basis of their respective movements. For example, people may time the imagined movement of the water so it corresponds with their perception of how fast or hard their wrist is working (Schwartz, in press). To affect this temporal coordination, they do not have to know the glass angle, only its rate of change. (As an analogy, one can monitor the speed of a clock's hands without knowing whether the clock is upright or sideways.) By this account, people solve the problem without detailed regard for how the glass angle would look if it were actually viewed. All they envision is the water rising along the side of the glass until it reaches the rim.

Consider the underestimated tilt of the thin cylinder in Experiment 1. Although there are many possible explanations, there are two theoretically relevant alternatives. One alternative is that people saw the angle of the glass in their mind's eye and were satisfied. Another alternative is that people were not aware of how the glass would look if visually realized. The current study tested these two alternatives. People tilted a glass with their eyes closed. Afterwards, they opened their eyes and made any adjustments they felt were necessary to get closer to the correct answer. If people do not adjust the glass, their image corresponds to their subsequent visual perception, and they are satisfied with the angle of their solution. But, if people do adjust the glass, their image is not like its equivalent visual perception.

To implement a suitable control, there was also a targeting task that required people to monitor the angle of the glass in their imagination. Individuals closed their eyes and tried to tilt a glass so that it pointed to 60° from upright. (They were told "10 o'clock" if right-handed or "2 o'clock" if left-handed.) Afterwards, they opened their eyes and were allowed to adjust the glass. Because this targeting task requires attention to the angular appearance of the glass, we expected that people would recode the wrist and glass behavior into a visual image. Consequently, when they open their eyes, they should make relatively few adjustments to the glass because they had imagined the angle. In contrast, for the pouring task, we thought people would increase their angle of tilt when they opened their eyes; they would discover the angle of the glass and its underestimation.

Method

Eight men and 8 women from the community completed both the pour and target conditions, with order of completion split in half and counterbalanced by gender. In each condition, individuals handled the thin glass with the 4.5-cm water level from Experiment 3. In the pour condition, they were given the usual instructions to close their eyes and tilt the glass until the water just reached the rim in their imaginations. Afterward, they opened their eyes and were given an opportunity to adjust their tilt if they wanted. With vision, they were allowed to tilt the glass forward or backward to make the adjustment (they could not start over). In the target condition, the same participants closed their eyes and were asked to tilt the glass so it pointed to 10 o'clock in the picture plane (2 o'clock if they were left-handed). Afterwards, they opened their eyes and were given the same opportunity to adjust the angle of the glass. We used clock face instructions because we thought they would be more communicative than "60°." We chose the target angle of 60° as a compromise that approximated the correct answer for the pour condition (73°) but still allowed for easy verbal communication.

Results

In the pour condition, 15 of the 16 participants increased the angle of the glass after they opened their eyes (the 16th made no adjustment). All 15 of the individuals increased the angle in the pour condition more than they did in the target condition. In the target condition, 5 individuals adjusted the glass forward, 1 backward, and 10 made no changes. Reflecting the different levels of adjustment, the correlation between the angle of the first tilt and the adjusted tilt was less in the pour condition than in the target condition (r^2s = .27 and .71, respectively). Figure 10 shows the closed-eye tilts for each condition and the subsequent with-vision adjustments. The Task × Adjustment interaction was significant, $F(1, 15) = 47.9$, $MSE = 27.2$. These results indicate that in the pour condition individuals systematically detected something with their eyes open that they were not aware of when their eyes were closed.

People's adjustments in the pour condition brought their tilts close to the correct angle. This indicates that the previous underestimated tilts for the thin glass were not intentional in the sense that people would have always chosen the under tilt. Performance in the target condition was approximately 15° short of the 60° target, even with vision. This is not overly inaccurate when one considers that 15° is half an hour on a clock face and individuals were trying to align the unmarked axis of a cylinder to the correct angle. It is not clear, however, why people underestimated, rather than overestimated, the angle. It may be that 45° (the average target response) is a special angle in wrist movement because elsewhere we have found that people can target 45° with great precision (Schwartz, in press).

Discussion

In the target condition, relatively few people changed the angle of the glass when given an opportunity to look at what they had done with their eyes closed. If we assume that they were imagining the glass, this result indicates the similarity of their imagery and its visual realization. In contrast, in the pour condition, nearly all of the individuals adjusted the glass when they opened their eyes. This indicates less of a similarity between people's imagery and subsequent vision. Our interpretation is that in the former case, people were

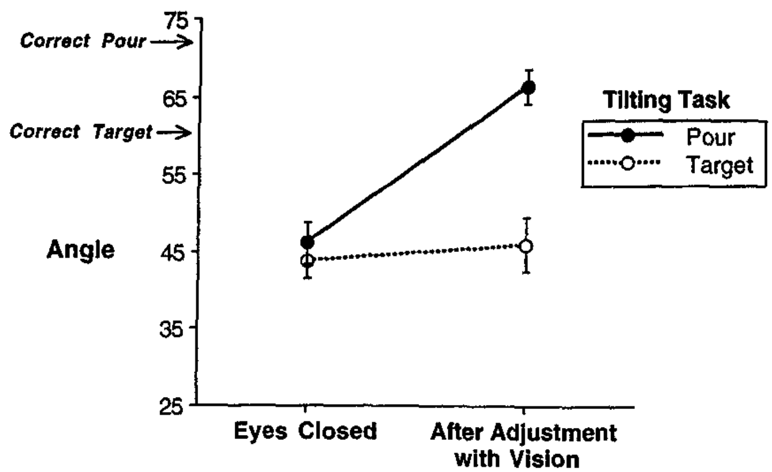

Figure 10. The adjustments that people made to their initial tilts for the pouring and targeting tasks after they opened their eyes (Experiment 5). Error bars represent the standard error.

explicitly attending to the angle per instructions, and therefore, the angle was a part of their image and corresponded to subsequent vision. In contrast, during the imagined pouring people did not visually imagine the angle of the glass nor were they aware of it. When they opened their eyes and explicitly saw the glass angle, they took this information into account and realized they had underestimated the tilt. So, although the angle of the glass was the outcome of importance for this task, people were not visually aware of it. In the General Discussion, we describe additional evidence to motivate our claim that simulated doing cannot be subsumed under the heading of visual imagery.

General Discussion

Empirical Summary

Five studies generated evidence to support the proposal that there is a form of problem solving that may be called knowing by simulated doing. In knowing by simulated doing, an (imagined) action not only facilitates an inference, it is the inference—independent of any visual awareness or descriptive evaluation of its truthfulness. It is a pragmatic form of knowledge in the extreme. There were three types of evidence: evidence that simulations are dissociable from descriptions, evidence that identifies representational properties of simulations, and evidence that simulations cannot be subsumed by visual imagery. We begin with a review of the first two types of evidence, and then we turn to the more theoretical task of distinguishing simulations from what has typically been studied under the heading of visual imagery. Afterwards, we consider implications of the current work with regard to naive physics.

The most straightforward contribution of the current research is evidence that people can solve a problem about a physical event through imagined actions, even though they do not know the answer descriptively. In Experiment 1, participants closed their eyes and tilted a thin and wide glass until some imagined water reached the rim. Almost everyone correctly tilted the thin glass farther than the wide one. Because this accuracy occurred for three different shapes, some unusual, it seems unlikely that people were simply replaying a memory. Instead, they were inferring the amount of tilt through their simulation of the water's behavior. In contrast, when people explicitly judged whether the pair of glasses would pour at the same angle, their inferences were rarely correct. Moreover, their tilting performance was not predictive of their accuracy in the judgment task. Experiment 2 further supported a dissociation of imagined actions and descriptions by showing that people did not have a propositional knowledge base, explicit or implicit, that could have led to the accurate tilting. People repeatedly switched among problem attributes to help them decide where to draw the water level in a target glass so that it would pour at the same angle as a standard glass. Yet, no matter what problem attribute they emphasized, people made the exact same error when they drew their conclusion for a given problem; they indicated that a wide glass needs a higher water level than a thin glass to tilt at the same angle (the opposite of the true answer). The inaccuracy of both the verbal and drawn judgments, coupled with the instability of belief, makes it hard to conceptualize how a tacit description could have led to the accurate tilting.

Experiments 1 and 2 showed that people could not solve the water pouring problem through description. The remaining experiments explored the properties of the imagined activity that enabled people to make the correct inference. Experiment 3 showed that people could complete the simulation solely in their imaginations and without overt motor activity. Experiment 3 also indicated that the simulation includes metric information because people's imagined tilts were sensitive to the water level. Experiment 4 showed that the simulations use causal propagation because people had difficulty going backward and against gravity. Finally, Experiment 5 indicated that people's simulations do not necessarily correspond to their visual perceptions.

Visual Imagery Versus Simulations of Doing

A focus of the current research has been the attempt to determine whether there are psychological components of simulated doing that are distinct from what has been typically explained under the umbrella of visual imagery. Although we believe simulated doing partakes of visual imagery, we do not believe it can be reduced to visual imagery. There are two reasons that we consider. The first has to do with the representation of physical forces and constraints, and the second has to do with the distinction between doing and viewing.

Forces and Constraints

Theoretical models of imagery typically concern the spatial information available to vision but not the force-based, physical information available to the haptic system (however, see Hubbard, 1995b). Kosslyn's (1980) seminal model of imagery, for example, used a coordinate system to represent position, configuration, and motion, but it had no way to represent weight or momentum. If people represent physical information to complete the water pouring task, then it suggests that knowing by simulated doing cannot be subsumed by traditional theories of visual imagery, and as described next, it counters theorizing that people cannot represent physical forces and constraints in analog form.

Shepard's (1984, 1994) account of analog imagery, as instantiated in his theory of resonant kinematics, explicitly rejected the representation of inanimate physical forces and constraints in analog form. This rejection comes in part from empirical evidence that imagery does not model Newtonian laws (e.g., Bundensen, Larsen, & Farrell, 1983; Gilden, 1991; Proffitt, Kaiser, & Whelan, 1990). There is, however, evidence of memory distortions that suggest analog representations of physical properties (e.g., Freyd, 1987; Freyd, Pantzer, & Cheng, 1988; Hubbard, 1995a, 1995b). And, there is perceptual and imagery-based evidence that people represent biomechanical constraints in analog form (e.g., Parsons, 1994; Shiffrar & Freyd, 1990; Viviani & Stucchi, 1992). Shepard's more theoretical reason for rejecting analog representations of physical constraints has to do with whether it is reasonable to assume that people evolved a set of analog representations and transformations for reasoning about physical input. For spatial relations, it might be adaptive to have evolutionarily fixed representations of certain invariant relationships (cf. Spelke, Breinlinger, Macomber, & Jacobson, 1992). For example, occlusion is an invariant correlate of depth. So, to ensure people's ability to anticipate depth relations, natural selection may have led to an imagery format with which people necessarily infer that an occluded object is farther away than an occluding object. (People cannot visualize the alternative.) The same logic cannot be applied to physical invariants because their effects are usually mediated by situation specific details. Even for a physical property as invariant and universal as gravity, it would be maladaptive to have evolved representational structures that necessarily enforce "gravity analogs" in the imagination. People would have no choice but to infer that a balloon would fall.

Even though questions about innate constraints on imagery are important, they do not seem directly relevant to the question of whether people have analog representations of physical information as we are proposing for simulated doing. Perhaps people can use learned, analog representations of forces and constraints on a situation-specific basis. For a simulation of one situation they might retrieve one analog constraint they had learned (e.g., sticky), and for another they might retrieve a different constraint (e.g., fluffy). We see no theoretical reason why this proposal is incompatible with prior evidence on imagery. Unfortunately, it is difficult to prove that people are reasoning with analog physical representations if those representations come and go depending on the circumstances and people's beliefs about those circumstances; there is always the chance that apparent analog physical reasoning has been functionally orchestrated by a descriptive (propositional) representation. This is one reason why the current results are important. They are an existence proof that people can intentionally simulate a distal physical behavior and that tacit descriptions are not surreptitiously responsible for the analog qualities of the simulation.

A critical question, however, is whether people actually represented an analog of physical causes to solve the water pouring task. The fact that people make an inference about a physical problem using imagery does not entail that those people are representing forces or physical constraints in analog form. The visual imagery alternative is that people solved the problem solely on the basis of spatial relations. Leyton (1989), for example, showed that people can infer the causal sequence that led to an object's deformed appearance on the basis of a few geometric principles. If people solved the current problem on the basis of spatial relations alone, then the current results would fit a visual model of imagery more comfortably.

One relevant piece of information comes from Experiment 4. In this experiment with the upside-down glasses, people had trouble solving the problem going backwards and against gravity, and they tended to convert the task into a problem about a forward flow of events (e.g., by imagining ice that thawed). If people were solving the problem solely on the basis of spatial relations, then one would not expect the change of direction to have interfered with their problem solving any more than asking people to rotate a letter clockwise or counterclockwise would interfere with a same-or-mirrored judgment. A second piece of relevant evidence comes from a set of studies that changed physical properties of the pouring problem without changing spatial ones (Schwartz, in press). In one study, for example, individuals completed the pouring task by imagining water or molasses. People tilted the imagined molasses farther than the imagined water, even though the spatial configurations of the two problems are the same. Interestingly, when asked afterwards, people explicitly stated that a glass of water and molasses would need to be tilted to the same angle. These results show that imagined physical constraints do affect people's simulations even when they are not explicitly aware of their effects.

Doing Versus Viewing

Theories of visual imagery are often about the similarity between visual experience and its representation. They are not about how people directly control the physical world to reach instrumental and informational goals. For theories of simulated viewing, it is not important to represent physical causality because the theories are not explaining how people change the world. But, for a theory of simulated doing, the representation of physical causality is fundamental. This is because "doing" requires taking advantage of causal forces and constraints to manipulate the world. Our assumption is that people need to have representations of how their embodied ideas will cause physical changes if they are to achieve their goals.

There are empirical implications that come with the emphasis on doing. In addition to the representation of physical causes, an implication developed earlier was the idea that people can complete successful actions without being visually aware of the "objective" variable they are controlling. Research on a simple motor task supports the claim (Bridgeman, 1992; Goodale, 1988). Individuals fixated on a dot in front of them. They were told that a second dot would appear to the side. Their task was to look and point at this second dot as soon as it appeared. As they shifted their eyes to look at the second dot, their eyes made a saccadic movement. At the peak velocity of the saccadic movement, a computer displaced the second dot. The participants were unaware of this movement. When asked afterwards, for example, participants did not know the dot had moved. The visual system does not yield perceptual awareness during a saccade. But, even though the participants did not see the dot move, their pointing movement landed directly on the dot after it was displaced.

For the case of physical inferences, our assumption is that simulated doing includes motor as well as visual components, and that the motor component may not always make its effects available to visual awareness or consciousness. Experiment 5 provided support for this claim. The results showed that when people opened their eyes after the pouring task, they consistently adjusted the angle of the glass to compensate for their original underestimation of the tilt. However, when people were asked to tilt a glass to a specific target angle, they did not adjust the glass after they opened their eyes. Thus, in the targeting case, people's imagery matched its visual realization, whereas in the pouring case, people's imagery did not match its visual realization. Our interpretation of this latter result is that people were not aware of the glass angle during their simulation; angle only became a factor when it was made salient through perception.

Elsewhere, we further showed that people do not solve the problem on the basis of how the angle of the glass looks (Schwartz, in press). Instead, people coordinate the imagined water movement with the rate of work they apply to the glass. Presumably, the motor system generates temporal information about the rate of work and this drives the movement of a timing-responsive water representation. In one study, weights were attached to the bottoms of glasses. In a targeting condition, people tilted a weighted and unweighted glass to 45°. People were quite accurate, with no differences between the weighted and unweighted glasses. Thus, weights do not cause people to misrepresent glass angle. Yet, when people completed the water pouring task, they tilted the weighted glass less than the unweighted one. Our interpretation is that people felt they were doing more work with the weighted glass. This increased the rate they transformed the imagined water, and the water reached the rim sooner. The difference between the targeting and pouring outcomes does not mean that people do not use visual imagery during the pouring task; rather it suggests that it is worthwhile distinguishing functional complexes of imagery. Whereas targeting relies on spatial imagery, physical inferences may rely on physical imagery (Schwartz, in press), or as we have called it here because we are emphasizing its basis in action, simulated doing.

Knowing by Simulated Doing and Naive Physics

The current experiments have not been intended to show that simulations are better than descriptions, nor that the many demonstrations of naive physics errors would all disappear if people could only make actions. Elsewhere, it has been argued that simulations and descriptions each have their own strengths, appropriate uses, and pathways of interaction (Schwartz, in press; Schwartz & Black, 1996b; Schwartz & Hegarty, 1996). The goal of the current research has simply been to explore whether people can extend their bodily methods of interacting with the physical world to solve problems about relatively distal physical behaviors. To make sure this is what we were studying, it was necessary to dissociate simulated doing from descriptive problem-solving and visual imagery. Next, we consider the prospects for putting these inferential methods back together again.

diSessa (1983) proposed that there is a type of knowledge, the phenomenological primitive, that underlies people's physical intuitions about the mechanisms of the world. These primitives, like springiness, "stand without significant explanatory substructure or justification" (p. 15). People's basic mechanical intuitions are not easily opened to reflection or decomposition. People, for example, know what a spring does but they cannot explain it much beyond saying it is "springy." If one assumes that the water pouring task relates to such primitives, the current results support some speculations on their obscurity. First, a primitive may often only manifest its inferences during (imagined) actions (e.g., when one imagines stretching a spring). This means that the primitive's inferential structure is timing based and transitory; it cannot be paused and reflected upon. Second, physical primitives may be bound up with the motor system. The "first person" phenomenology of motor representations (e.g., of force) may be difficult to re-represent visually and make available for sustained reflection. Finally, as we describe below, people's explicit descriptions can commandeer their simulations so the simulations try to portray those descriptions. For example, people can imagine themselves flapping their wings and flying. In this case, description drives visual imagery that in turn overwhelms or "contaminates" (Wong & Mack, 1981) simulated doing. As a consequence, people lose access to their primitives. This last possibility has important implications for instruction and theories of embodied cognition.

Although the primitives that operate through simulated doing are hard to inspect, they probably ground more abstract descriptions (e.g., Clement, 1994; diSessa, 1993). The notion of force, for example, probably comes through the haptic system. There is, however, relatively little research about how nonverbal primitives become integrated into the more descriptive knowledge of scientific discourse (Alibali & Goldin-Meadow, 1993; Intons-Peterson, 1993; Tversky & Schiano, 1989). This is an important issue to investigate because the answer may be educationally relevant and because there is a common assumption that embodied cognition grounds higher order cognition (e.g., motor imagery; Grush, 1994). The current research provides some warnings on this issue.

Consider the fact that people often underestimate the tilt of the glass when their eyes are closed. How may we improve this performance? This may be a difficult task. By our hypothesis, people do not conduct their simulations by attending to the objective control variable of glass angle. Consequently, if they are told to adjust their behavior according to the objective control variable, it may block or distort the variables that make their intuitive simulations reasonably effective. Imagine, for example, that people do not begin by tilting the glasses as in the current experiments, but instead, they are first asked to judge explicitly whether the wide and thin glasses pour at the same angle. This would focus their attention on the angle of the glass. What would this attention to the angle, the correct control variable, do to a subsequent simulation of the wide and thin glass? Assuming that people do not normally attend to the angle in their simulations, the attention to the angle should contaminate their simulations.

A test of this hypothesis confirmed the prediction. Not only did participants get their simulations wrong, but they did not even make the simulations fit their erroneous judgments (Schwartz, in press). When people were encouraged to use the glass angle to regulate a simulation, it undermined their ability to harness their physical intuitions into an effective simulation. Perhaps it would be more effective to arrange for feedback that helped people discern the variables that their actions control. Regardless, this result and the current experiments make it clear that the intuitive knowledge found in simulated doing is not always on the same plane as other forms of knowledge. As a consequence, one cannot simply assume that the primitive and embodied knowledge of (imagined) action naturally evolves into more reflective forms of understanding. It is not a trivial matter to connect epistemic planes that draw inferences using fundamentally different variables. In this light, one may understand why naive physical intuition and formal instruction do not always make contact.

References

Alibali, M. W., & Goldin-Meadow, S. (1993). Gesture–speech mismatch and mechanisms of learning: What the hands reveal about a child's state of mind. *Cognitive Psychology, 25,* 468–523.

Anderson, J. R. (1978). Arguments concerning representations for mental imagery. *Psychological Review, 85,* 249–277.

Black, T., & Schwartz, D. L. (1996, June). *When imagined actions speak louder than words.* Paper presented at the annual meeting of the Jean Piaget Society, Philadelphia.

Bridgeman, B. (1992). Conscious vs. unconscious processes: The case of vision. *Theory & Psychology, 2,* 73–88.

Brown, J. S., & Burton, R. R. (1977). *Diagnostic models for procedural bugs in basic mathematical skills.* Cambridge, MA: Bolt, Bernaek & Newman.

Bundensen, C., Larsen, A., & Farrell, J. E. (1983). Visual apparent movement: Transformations of size and orientation. *Perception, 12,* 549–558.

Clement, J. (1983). A conceptual model discussed by Galileo and used intuitively by physics students. In D. Gentner & A. L. Stevens (Eds.), *Mental models* (pp. 325–340). Hillsdale, NJ: Erlbaum.

Clement, J. (1994). Use of physical intuition and imagistic simulation in expert problem solving. In D. Tirosh (Ed.), *Implicit and explicit knowledge* (pp. 204–242). Hillsdale, NJ: Erlbaum.

Cooke, N. J., & Breedin, S. D. (1994). Constructing naive theories of motion on the fly. *Memory & Cognition, 22,* 474–493.

Decety, J., & Michel, F. (1989). Comparative analysis of actual and mental movement times in two graphic tasks. *Brain and Cognition, 11,* 87–97.

de Kleer, J., & Brown, J. S. (1984). A qualitative physics based on confluences. *Artificial Intelligence, 24,* 7–83.

diSessa, A. A. (1983). Phenomenology and the evolution of intuition. In D. Gentner & A. L. Gentner (Eds.), *Mental models* (pp. 15–33). Hillsdale, NJ: Erlbaum.

diSessa, A. A. (1993). Toward an epistemology of physics. *Cognition and instruction, 10,* 105–225.

Finke, R. A. (1980). Levels of equivalence in imagery and perception. *Psychological Review, 87,* 113–132.

Finke, R. A. (1985). Theories relating mental imagery to perception. *Psychological Bulletin, 98,* 236–259.

Finke, R. A., & Freyd, J. J. (1989). Mental extrapolation and cognitive penetrability: Reply to Ranney and proposals for evaluative criteria. *Journal of Experimental Psychology: General, 188,* 403–408.

Forbus, K. D. (1988). Qualitative physics: Past, present and future. In H. Shrobe (Ed.), *Exploring artificial intelligence* (pp. 239–296). San Mateo, CA: Morgan Kaufman.

Freyd, J. J. (1983). The mental representation of movement when viewing static stimuli. *Perception & Psychophysics, 33,* 575–581.

Freyd, J. J. (1987). Dynamic mental representations. *Psychological Review, 94,* 427–438.

Freyd, J. J., Pantzer, T. M., & Cheng, J. L. (1988). Representing statics as forces in equilibrium. *Journal of Experimental Psychology: General, 117,* 395–407.

Gardin, F., & Meltzer, B. (1989). Analogical representations of naive physics. *Artificial Intelligence, 38,* 138–159.

Gentner, D., Brem, S., Ferguson, R. W., Markman, A. B., Levidow, B. B., Wolff, P., & Forbus, K. (1997). Analogical reasoning and conceptual change: A case study of Johannes Kepler. *Journal of the Learning Sciences, 6,* 3–40.

Gibson, J. J. (1962). Observations on active touch. *Psychological Review, 69,* 477–491.

Gilden, D. (1991). On the origins of dynamical awareness. *Psychological Review, 98,* 554–568.
Goodale, M. A. (1988). Modularity in visuomotor control: From input to output. In Z. W. Pylyshyn (Ed.), *Computational processes in human vision. An interdisciplinary perspective* (pp. 262–285). Norwood, NJ: Ablex.
Grush, R. (1994). Motor models as a step to higher cognition. *Behavioral and Brain Sciences, 17,* 209–210.
Hecht, H. (1996). Heuristics and invariants in dynamic event perception: Immunized concepts or nonstatements? *Psychonomic Bulletin & Review, 3,* 61–70.
Hegarty, M. (1992). Mental animation: Inferring motion from static displays of mechanical systems. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 18,* 1084–1102.
Hegarty, M., Just, M. A., & Morrison, I. R. (1988). Mental models of mechanical systems: Individual differences in qualitative and quantitative reasoning. *Cognitive Psychology, 20,* 191–236.
Hubbard, T. L. (1995a). Cognitive representations of motion: Evidence for friction and gravity analogues. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 21,* 241–254.
Hubbard, T. L. (1995b). Environmental invariants in the representation of motion: Implied dynamics and representational momentum, gravity, friction, and centripetal force. *Psychonomic Bulletin & Review, 2,* 322–338.
Hubbard, T. L. (1997). Target size and displacement along the axis of implied gravitational attraction: Effects of implied weight and evidence of representational gravity. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 23,* 1484–1493.
Intons-Peterson, M. J. (1993). Imaginal priming. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 19,* 223–235.
Jeannerod, M. (1994). The representing brain: Neural correlates of motor intention and imagery. *Behavioral and Brain Sciences, 17,* 187–245.
Kaiser, M. K., Jonides, J., & Alexander, J. (1986). Intuitive reasoning about abstract and familiar physics problems. *Memory & Cognition, 14,* 308–312.
Kosslyn, S. M. (1976). Can imagery be distinguished from other forms of internal representation? Evidence from studies of information retrieval times. *Memory & Cognition, 4,* 291–297.
Kosslyn, S. M. (1980). *Image and mind.* Cambridge, MA: Harvard University Press.
Krist, H., Fieberg, E. L., & Wilkening, F. (1993). Intuitive physics in action and judgment: The development of knowledge about projectile motion. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 19,* 952–966.
Lackner, J. R. (1992). Spatial orientation in weightless environments. *Perception, 21,* 803–812.
Lederman, S. J., Klatzky, R. L., Chataway, C., & Summers, C. D. (1990). Visual mediation and the haptic recognition of two-dimensional pictures of common objects. *Perception & Psychophysics, 47,* 54–64.
Leyton, M. (1989). Inferring causal history from shape. *Cognitive Science, 13,* 357–387.
Liben, L. S., & Goldbeck, S. L. (1984). Performance on Piagetian horizontality and verticality tasks: Sex-related differences in knowledge of relevant physical phenomena. *Developmental Psychology, 20,* 595–606.
McAfee, E. A., & Proffitt, D. R. (1991). Understanding the surface orientation of liquids. *Cognitive Psychology, 23,* 483–514.
McCloskey, M. (1983). Naive theories of motion. In D. Gentner & A. L. Stevens (Eds.), *Mental models* (pp. 299–324). Hillsdale, NJ: Erlbaum.
McCloskey, M., Caramazza, A., & Green, B. (1980). Curvilinear motion in the absence of external forces: Naive beliefs about the motion of objects. *Science, 210,* 1139–1141.
Milner, A. D., & Goodale, M. A. (1995). *The visual brain in action.* Oxford, United Kingdom: Oxford University Press.
Parsons, L. M. (1994). Temporal and kinematic properties of motor behavior reflected in mentally simulated action. *Journal of Experimental Psychology: Human Perception and Performance, 20,* 709–730.
Pascual-Leone, J., & Morra, S. (1991). Horizontality of water level: A neo-Piagetian developmental review. *Advances in Child Development and Behavior, 23,* 231–273.
Pelisson, D., Prablanc, C., Goodale, M. A., & Jeannerod, M. (1986). Visual control of reaching movements without vision of the limb. II. Evidence of fast unconscious processes correcting the trajectory of the hand to the final position of a double-step stimulus. *Experimental Brain Research, 62,* 303–311.
Perky, C. W. (1910). An experimental study of the imagination. *American Journal of Psychology, 21,* 422–452.
Posner, M. I. (1967). Characteristics of visual and kinaesthetic memory codes. *Journal of Experimental Psychology, 75,* 103–107.
Proffitt, D. R., Kaiser, M. K., & Whelan, S. M. (1990). Understanding wheel dynamics. *Cognitive Psychology, 22,* 342–373.
Pylyshyn, Z. W. (1973). What the mind's eye tells the mind's brain: A critique of mental imagery. *Psychological Bulletin, 80,* 1–24.
Pylyshyn, Z. W. (1981). The imagery debate: Analog media versus tacit knowledge. In N. Block (Ed.), *Imagery* (pp. 151–206). Cambridge, MA: MIT Press.
Rieser, J. J. (1990). Development of perceptual-motor control while walking without vision: The calibration of perception and action. In H. Bloch & B. I. Bertenthal (Eds.), *Sensory-motor organizations and development in infancy and early childhood* (pp. 379–408). Dordrecht, the Netherlands: Kluwer.
Rizzolatti, G., Camarda, R., Fogassi, L., Genilucci, M., Luppino, G., & Matelli, M. (1988). Functional organization of area 6 in the macaque monkey. II. Area F5 and the control of distal movements. *Experimental Brain Research, 71,* 491–507.
Rock, I., Wheeler, D., & Tudor, L. (1989). Can we imagine how objects look from other viewpoints? *Cognitive Psychology, 21,* 185–210.
Schooler, J. W., Ohlsson, S., & Brooks, K. (1993). Thoughts beyond words: When language overshadows insight. *Journal of Experimental Psychology: General, 122,* 166–183.
Schwartz, D. L. (in press). Physical imagery: Dynamic versus kinematic models. *Cognitive Psychology.*
Schwartz, D. L., & Black, J. B. (1996a). Analog imagery in mental model reasoning: Depictive models. *Cognitive Psychology, 30,* 154–219.
Schwartz, D. L., & Black, J. B. (1996b). Shuttling between depictive models and abstract rules: Induction and fallback. *Cognitive Science, 20,* 457–497.
Schwartz, D. L., & Hegarty, M. (1996). Coordinating multiple representations for reasoning about mechanical devices. In P. Olivier (Ed.), *AAAI Spring Symposium: Cognitive and computational models of spatial representation* (pp. 101–109). Menlo Park, CA: AAAI Press.
Shepard, R. N. (1984). Ecological constraints on internal representation: Resonant kinematics of perceiving, imagining, thinking, and dreaming. *Psychological Review, 91,* 417–447.
Shepard, R. N. (1994). Perceptual-cognitive universals as reflections of the world. *Psychonomic Bulletin & Review, 1,* 2–28.
Shepard, R. N., & Cooper, L. A. (Eds.). (1986). *Mental images and their transformation.* Cambridge, MA: MIT Press.
Shiffrar, M., & Freyd, J. J. (1990). Apparent motion of the human body. *Psychological Science, 1,* 257–264.

Spelke, E. S., Breinlinger, K., Macomber, J., & Jacobson, K. (1992). Origins of knowledge. *Psychological Review, 99,* 605–632.

Turvey, M. T. (1996). Dynamic touch. *American Psychologist, 51,* 1134–1152.

Turvey, M. T., Solomon, H. Y., & Burton, G. (1989). An ecological analysis of knowing by wielding. *Journal of the Experimental Analysis of Behavior, 52,* 387–407.

Tversky, B., & Schiano, D. J. (1989). Perceptual and conceptual factors in distortions in memory for graphs and maps. *Journal of Experimental Psychology: General, 118,* 387–398.

Viviani, P., & Stucchi, N. (1992). Biological movements look uniform: Evidence of motor-perceptual interactions. *Journal of Experimental Psychology: Human Perception and Performance, 18,* 603–623.

Vosniadou, S., & Brewer, W. F. (1993). Mental models of the earth: A study of conceptual change in childhood. *Cognitive Psychology, 24,* 535–585.

White, B. (1993). ThinkerTools: Causal models, conceptual change and science education. *Cognition and Instruction, 10,* 1–100.

White, B., & Frederiksen, J. (1990). Causal model progressions as a foundation for intelligent learning environments. *Artificial Intelligence, 24,* 99–157.

Williams, M. D., Hollan, J. D., & Stevens, A. L. (1983). Human reasoning about a simple physical system. In D. Gentner & A. L. Stevens (Eds.), *Mental models* (pp. 131–153). Hillsdale, NJ: Erlbaum.

Wohlschläger, A., & Wolschläger, A. (1998). Mental and manual rotation. *Journal of Experimental Psychology: Human Perception and Performance, 24,* 397–412.

Wong, E., & Mack, A. (1981). Saccadic programming and perceived location. *Acta Psychologica, 48,* 123–131.

Received December 8, 1997
Revision received June 2, 1998
Accepted June 3, 1998 ■