



Sources of Uncertainty in Intuitive Physics

Kevin A. Smith, Edward Vul

Department of Psychology, University of California, San Diego

Received 19 September 2012; accepted 26 September 2012

Abstract

Recent work suggests that people predict how objects interact in a manner consistent with Newtonian physics, but with additional uncertainty. However, the sources of uncertainty have not been examined. In this study, we measure perceptual noise in initial conditions and stochasticity in the physical model used to make predictions. Participants predicted the trajectory of a moving object through occluded motion and bounces, and we compared their behavior to an ideal observer model. We found that human judgments cannot be captured by simple heuristics and must incorporate noisy dynamics. Moreover, these judgments are biased consistently with a prior expectation on object destinations, suggesting that people use simple expectations about outcomes to compensate for uncertainty about their physical models.

Keywords: Intuitive physics; Stochastic simulation; Uncertainty; Probabilistic inference

1. Introduction

Predicting how the world will unfold is key to our survival and ability to function on a daily basis. When we throw a ball, cross a busy street, or catch a pen about to fall off of a desk, we must foresee the future physical state of the world to plan our actions. The cognitive mechanisms that help us make these predictions have been termed “intuitive physics” models.

Although human performance in physical prediction tasks tends to approximate real-world (Newtonian) physics, it does not match exactly: People make systematic prediction errors. While this has been taken as evidence that human models of intuitive physics are non-Newtonian (e.g., McCloskey, 1983), more recently human behavior has been explained by intuitive Newtonian physics models under uncertainty. On this account, human predictions deviate from Newtonian mechanics because of stochastic error—uncertainty about the initial positions or velocities of objects propagates through the non-linear

Correspondence should be sent to Kevin A. Smith, Department of Psychology, University of California, San Diego, 9500 Gilman Drive, La Jolla, CA 92093-0109. E-mail: k2smith@ucsd.edu

physical model and causes variability and bias in final judgments. For instance, human predictions about the stability of a tower of blocks or the most likely direction for that tower to fall are consistent with a purely Newtonian model of physics with a small amount of uncertainty in the initial positions of the constituent blocks (Hamrick, Battaglia & Tenenbaum, 2011). Similar models of physics with perceptual noise have been used to explain relative mass judgments in collisions (Sanborn, Mansinghka & Griffiths, 2009) and infants' expectations for object movement (Téglás et al., 2011).

There are numerous ways in which uncertainty can be introduced into intuitive physical reasoning. We broadly classify these into two categories: perceptual uncertainty and uncertainty about dynamics. Perceptual uncertainty arises because initial measurements of the location and velocity of objects is imperfect; this initial noise propagates through the model. Uncertainty about dynamics reflects noise in the physical model itself. Real object movement and collisions are perfectly deterministic only in an idealized system; in the world, objects can deviate from their ideal path because of multiple, unknowable interactions with the environment (e.g., a ball rolling across gravel will not move in a straight line). Stochastic dynamics could thus reflect such environmental uncertainty.

The goal of the study was to disentangle the influence of initial noisy percepts and noisy physics on human predictions of object dynamics. We compared human behavior in a simple physical prediction task to a stochastic physics model with parameters reflecting the different types of uncertainty.

2. Stochastic physics model

We designed a model to replicate stochastic physics in a simple environment: a ball bouncing around a two-dimensional box. We based this model on idealized mechanics but also incorporated the two sources of uncertainty: We added noise to the initial position and velocity to capture perceptual uncertainty, while dynamic uncertainty was captured by jitter in object movement over time, and variability in bounce angles.

2.1. Uncertainty parameters

The model was based on a simple two-dimensional physics engine customized to add our sources of uncertainty. As physical uncertainty goes to zero, this model reduces to laws from idealized mechanics: The ball would continue to move in a straight line at a constant velocity until it hit a wall, at which point it would bounce elastically and with angle of incidence equal to the angle of exit. Uncertainty was captured using four parameters, two for the perceptual error, and two for the stochastic error (see Fig. 1).

2.1.1. Perceptual uncertainty

At the start of the simulation, the ball's position and velocity were based on where the ball would be in a perfectly deterministic simulation, but with noise added. Position was perturbed by isotropic two-dimensional Gaussian noise parameterized by standard




Fig. 1. Sources of uncertainty in the stochastic physics model. Each parameter refers to a different source of noise: position noise (σ_p), velocity direction noise (κ_v), ongoing movement noise (κ_m), and noise added with a bounce (κ_b).

deviation, σ_p . Noise in velocity direction was captured in a von Mises (circular normal) distribution on direction of motion, parameterized by concentration (inverse variance) κ_v . We did not consider uncertainty in the speed of the ball, as this would only affect the timing of the ball's movement but not its destination, which is the prediction we aim to capture.

2.1.2. Dynamic uncertainty

Noise was added during the simulation in two ways. First, at each time step (1,000/s), the direction of the ball was "jittered" by adjusting its direction using a von Mises distribution with the concentration parameter κ_m . In addition, noise was added during each bounce by assuming that the angle the ball bounced off of the wall was defined by a von Mises distribution centered on the angle of incidence with a concentration parameter κ_b .

3. Experiment

We aimed to test model predictions against human data and to estimate uncertainty parameters in intuitive dynamics. In this experiment, subjects predicted the trajectory of a ball in a two-dimensional environment on a computer screen. This was performed in a "Pong" game where participants tried to catch the ball with a paddle. Crucially, we occluded the latter part of the ball's movement, so that successful prediction of the final position required the mental simulation of the object trajectory. We could estimate the final position predicted by our stochastic physics model with different parameters, and thus compare human behavior to model predictions under varying types and degrees of uncertainty.

In this experiment, we parametrically varied both the distance the ball would travel¹ and the number of bounces off of walls while occluded. If intuitive dynamics models are deterministic, then the number of bounces will have no effect on human predictions. The

distance manipulation was designed to tease apart the contributions of perceptual uncertainty about velocity and dynamic velocity noise.

3.1. Methods

Fifty-two UCSD undergraduates (with normal or corrected vision) participated in the experiment for course credit.

Subjects used a computer mouse to control the vertical position of an on-screen “paddle” to catch a moving ball. The ball moved according to the deterministic physics underlying the stochastic physics model. Both the paddle and the ball were confined to a 1200×900 pixel area in the center of the screen. Each trial began with a display of only the paddle, which subjects could move up and down. The paddle was 100 pixels in height and was centered on the vertical position of the mouse before each trial (Fig. 2a). A mouse click triggered the start of a trial. A ball would then appear on the screen, moving at a constant velocity of 600 pixels/s. After the ball moved 400 pixels (667 ms), a gray rectangle would occlude the portion of the screen containing the ball (Fig. 2b). During this period, the ball would continue to move, bouncing perfectly elastically off of the edges of the field, but would not be visible. Once the subjects caught the ball with the paddle, or the ball broke the plane of the paddle, the trial would end and the occluder would be removed, showing whether (and by how far) the subject missed the ball (Fig. 2c). Upon clicking the mouse, the screen would clear and reset for the next trial. The number of balls caught by the subject was always displayed in the upper right corner as a motivation to perform well.

Subjects were given 648 trials throughout the experiment. These 648 trials were identical for all subjects but were presented in a randomized order. Each trial had a particular ball trajectory, generated by one of nine conditions. The nine trajectory conditions crossed the distance the ball travelled while occluded (600, 800, or 1,000 pixels) with the number of bounces (0, 1, or 2); there were 72 trials of each condition. The specific path for each trial was generated prior to the experiment subject to the constraints of the condition and the constraint that the final position was not in the top 20% or bottom 20% of the enclosed area to avoid bias due to positioning the paddle at the ends of the screen.




Fig. 2. Diagram of a trial. (A) The ball moves unoccluded in a straight line. (B) Once the field is occluded, the ball continues to move and the subject must predict where it will end. (C) The trial ends once the ball is either caught or passes the plane of the paddle.

Before starting the experiment, subjects were given seven trials without the occluder to demonstrate how the ball would move, then six practice trials with the occluder.

For each trial, we recorded the position of the midpoint of the paddle once the ball was caught or moved past the paddle. From this measure, we could calculate, for each trial: (a) the average expected position of the ball and (b) the variance of predictions around that expectation.

3.2. Subject performance

3.2.1. Accuracy

Subjects caught the ball on 43.8% of all trials. Individual subject accuracies varied between 25.6% and 63.7% (chance was 11%). Accuracy also varied by trial condition: Subjects were most accurate in the shortest, no bounce condition (69%) and least accurate in the longest, two-bounce condition (32%).

Accuracy improved slightly over time, increasing from 42.7% in the first half of trials to 44.9% on the second half ($\chi^2(1) = 15.9, p < .001$). However, because this was a small effect, and because in a logistic model predicting accuracy, trial order did not interact with either distance ($\chi^2(2) = 0.72, p = .70$) or number of bounces ($\chi^2(2) = 4.18, p = .12$), we do not try to account for this change.

3.2.2. Expected positions

In addition to decreasing accuracy, subjects also showed increasing bias in average predictions as the distance or number of bounces increased. The mean final position of the paddle for each trial shifted toward the center as compared to the final ball position (see Fig. 3). The magnitude of this bias toward the center of the screen increased as either distance or number of bounces increased (Table 1).

3.2.3. Variance of responses

The variability of subjects' responses around the mean also increased with distance and bounces, but only up to a ceiling—well below the maximum possible spread—once subjects had to take into account even one bounce (Table 2).

4. Model application

The coarse results suggest that prediction error and variability increase with distance or number of bounces. But they do not indicate which sources of uncertainty contribute to intuitive physics predictions, nor do they explain why some trials within the same condition produce greater bias and variability than others.

We aimed to tease these factors apart via our model of stochastic physics. By finding the set of uncertainty parameters that best fits the empirical data, we can compare the relative contribution of the perceptual uncertainty parameters to the dynamic uncertainty

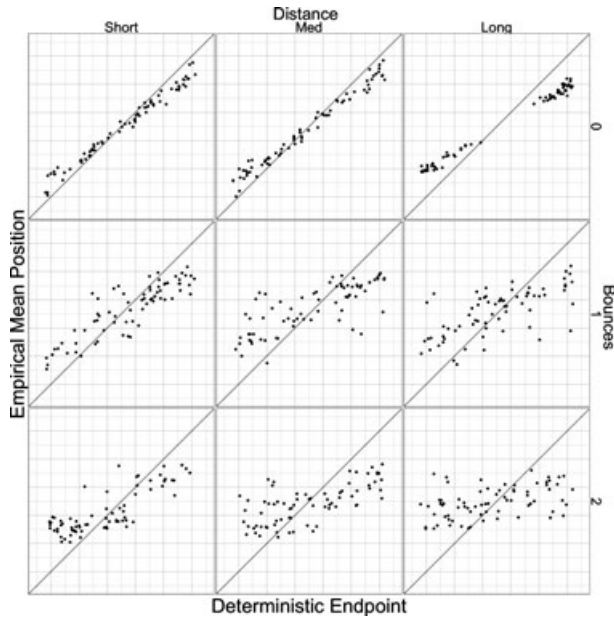


Fig. 3. Mean predicted paddle position versus path endpoint using deterministic physics as a function of trial condition. Each point represents a separate trial.

Table 1
Percent of distance “shifted” from actual end ball position toward center by trial condition

	Distance		
	600	800	1000
Bounces			
0	24%	44%	53%
1	23%	60%	70%
2	41%	63%	84%

Table 2
Average standard deviation (in pixels) of responses within a trial by condition

	Distance		
	600	800	1000
Bounces			
0	65	76	94
1	111	115	114
2	115	111	121

parameters. A good model should capture trial-level differences in subjects' performance and explain trial difficulty based on the interplay of different sources of uncertainty.

4.1. Simulation

We replicated the experimental task in the stochastic physics model, simulating the same 648 trials. To mirror this task, each simulation started at the point of occlusion (when subjects could no longer visually track the ball and must predict its path) and ended when the simulated ball crossed the plane of the paddle. On each simulation, we measured the position of the simulated ball along that plane. Because there is no analytic form of the probability distribution over possible trajectories, we simulated each trial 500 times, thus estimating the predictive distribution for each trial via sampling.

No reasonable set of uncertainty parameters produced mean estimates of the final position of the ball that were systematically shifted toward the center like the empirical data; as long as Newtonian physics underlies the model, averaging over all simulation paths, the mean ending position will be close to the actual endpoint for most trials, regardless of the uncertainty parameters chosen.² As the magnitude of the center bias scaled with distance and number of bounces, we suspected that subjects were incorporating a prior on final position, producing a center bias proportional to the uncertainty in their physics-based predictions. People therefore appear to incorporate prior expectations with their intuitive physics models.

We treated this bias as a simple Gaussian prior on the final ball position centered on the middle of the screen, with standard deviation as a free parameter (σ_{prior}). One value of this parameter was used for all trials and conditions.

The final distribution of predictions for each trial was calculated by combining the center-prior with the distribution of predicted positions simulated by the stochastic physics engine. We treated the distribution of predicted positions as a Gaussian and calculated their mean and standard deviation. We could then calculate the mean and standard deviation of the posterior distribution using Bayesian cue combination (e.g., Ernst & Banks, 2002):

$$\sigma_{\text{post}}^2 = \left(\frac{1}{\sigma_{\text{prior}}^2} + \frac{1}{\sigma_{\text{sim}}^2} \right)^{-1} \quad \mu_{\text{post}} = \left(\frac{x_{\text{center}}}{\sigma_{\text{prior}}^2} + \frac{\mu_{\text{sim}}}{\sigma_{\text{sim}}^2} \right) * \sigma_{\text{post}}^2$$

Using these equations, trials with greater simulation variance will be more affected by the prior and will shift further toward the screen center. Thus, the model can account for the center-bias in a manner sensitive to prediction uncertainty.

We found the maximum likelihood parameters to fit three quarters of the data (with an equal number of trials from each of the distance by bounce conditions).³ We also fit two other models: one with only perceptual uncertainty and prior parameters, and other with only dynamic uncertainty and prior parameters. We compared these models based on the likelihood of the quarter of the remaining (cross-validation) data.

4.2. Model results

4.2.1. Model comparison

We designed the stochastic physics model to investigate how various sources of uncertainty contribute to intuitive physics. Thus, we compared the model with both dynamic and perceptual uncertainty to the two nested models with either dynamic or perceptual uncertainty parameters alone to determine which sets of parameters were necessary to best explain the data (see Fig. 4).

In addition, we tested how well any of the stochastic models captures human behavior by comparing them to a “heuristic oracle” model with different parameters for each condition. The heuristic oracle model assumes that people know the correct answer (thus “oracle”) but produce errors that vary by condition without regard to individual trial details (“heuristic”). These errors include some bias toward the center (given by a linear relationship between average reported position and the deterministic end point) and response variability distributed around that shifted position (with variance estimated independently for each condition). In other words, the heuristic oracle model is a non-physical error model. This model can capture the gross “shift” in expected position that was observed in the data in each condition (see Fig. 3), but it does not treat the shift as an inference done independently on each trial. The spread in responses was assumed to be constant within each condition and was set at the average empirical standard deviation from that condition. Like the stochastic models, this model was fit on three-quarters of the trials and tested on the remaining data.

Table 3 shows cross-validation likelihood for the four models. All log-likelihoods are shown as improvement over a baseline assuming that all data came from a single Gaussian. In addition, we included a “perfect trial fit” model that knows the mean and standard deviation of responses for each trial—this serves as the plausible upper limit on how well different models might do. The full stochastic model does best, followed closely by a




Fig. 4. Sample simulation paths for one trial with each model. The gray lines represent individual simulations; the black line represents deterministic simulation. There is no initial uncertainty in the dynamic model, but it builds quickly over time, resulting in wavy paths. The initial position and velocity vary significantly in the perceptual model, but once started, the simulation unfolds deterministically. The full model uses both types of uncertainty and so has more certainty in starting positions than the perceptual model and straighter paths than the dynamic model.

Table 3
Model prediction of left-out data

Model	ΔLH
Full	2,588
Dynamic	2,568
Perceptual	2,197
Heuristic oracle	2,326
Perfect trial fit	3,259

model including only dynamic noise. Both the perceptual noise model and the non-physical model perform worse by many orders of magnitude.

The dynamic model performed nearly as well as the full model for two reasons. First, the parameter representing error in the initial position (σ_p) was set to a small value in the full model and explained very little of the variance in simulations. Second, much of the noise in initial velocity direction (κ_v) can be captured by increasing dynamic velocity noise (κ_m), and so we cannot say whether any initial velocity noise is required. The model with only perceptual noise did quite poorly because subjects' performance changed with each additional bounce, and thus human performance cannot be captured without dynamic uncertainty.

4.2.2. Trial-level simulations

Human predictions about individual trials within the same distance-by-bounce condition varied significantly: Some had much larger variations in responses or greater shifts toward the center than others. These differences arose from trajectory characteristics other than total distance traveled or number of bounces. For instance, it is harder to predict the end position of a ball that bounces in a corner or balls that approach the paddle at a steep angle. If the stochastic physics model is capturing characteristics of intuitive physics, then it should capture this within-condition variability as well.

The full stochastic model fit the variation in mean paddle position across trials well ($r = .93$), and slightly better than the predictions of the heuristic oracle model ($r = .90$). However, the difference between models is highlighted when considering individual conditions: Although both models account for the mean position in the no-bounce conditions, only the full model continues to perform well as bounces and distance are added (see Table 4).

Table 4
Correlation between model and empirical by-trial means within distance and bounce condition

	Full			Heuristic oracle		
	Distance			Distance		
	600	800	1000	600	800	1000
Bounces						
0	.99	.99	.99	.99	.99	.99
1	.86	.88	.85	.88	.77	.68
2	.89	.87	.82	.82	.68	.45




Fig. 5. Full model versus empirical standard deviation by trial. Each point represents a separate trial.

The standard deviation of predictions from the full stochastic model was well correlated with the standard deviation of subjects’ responses across trials ($r = .79$, see Fig. 5), albeit with a tendency to overestimate. Moreover, the stochastic physical model also captures the variability across trials within each distance-by-bounce condition (Table 5). Together, these results indicate that human uncertainty about final outcomes accumulates in a manner qualitatively similar to that predicted by a stochastic physical model.

In the experimental data, the amount of mean-shifting for each trial is related to the variance of the observations from that trial (Spearman’s rho = 0.30), suggesting that people hedge their guesses toward the middle more as the amount of uncertainty increases. A center-prior captures this behavior by causing more reliance on the prior when there is a wider distribution of model simulations. This has the effect of shifting guesses more toward the center when physical simulations are more uncertain. The stochastic physics model captures this phenomenon by predicting trial-level differences in uncertainty and is

Table 5
Correlation between full model and empirical by-trial standard deviations within condition

	Distance		
	600	800	1,000
Bounces			
0	.54	.43	.17
1	.53	.44	.30
2	.14	.16	.17

thus better able to describe variation in human responses across trials than a constant mean-shift for each condition (see Fig. 6).

4.3. Source of the center bias

Subjects positioned their paddle closer to the middle of the screen than where the ball actually ended, and we suggest that this bias arises from subjects' prior expectation that the ball will end in the center. In this section we address alternate explanations for this bias: Is the center bias arising from task demands and strategies for dealing with this difficult task? Or is such a bias learned over time? We argue that neither of these accounts explains the bias we observe.

We assume that the middle of the paddle is each subject's best guess for the end position of the ball, but subjects could instead be attempting to minimize a loss function on the distance between each of the simulation outputs and where they place the paddle. Because predicted end-points under a physical model are somewhat skewed away from the edges (toward the center) due to the physical non-linearities of bounces, estimated positions will also be skewed toward the center relative to the modes of the distributions. However, these effects do not explain subjects' center shift. With a quadratic loss function (L2), the best placement of the paddle would be the mean of the simulations (Strook, 2011, p. 43), but as noted previously, the mean of the distributions was often centered on the end position of the ball and does not account for the observed center bias (indeed, this is why we suspected that

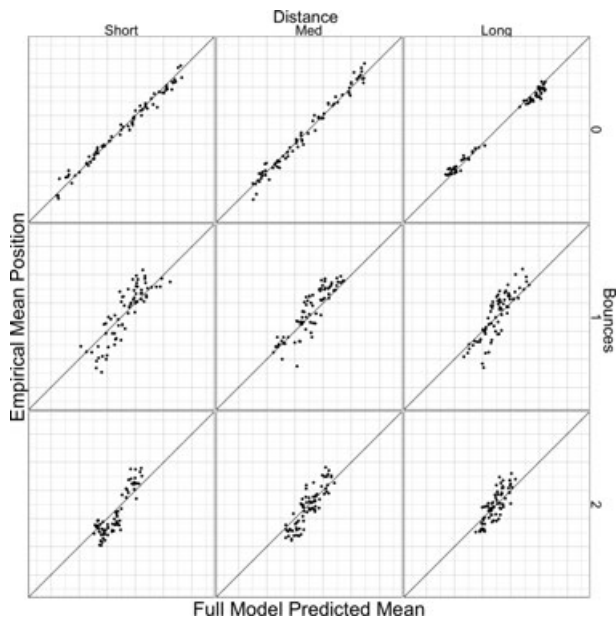


Fig. 6. Full model predictions versus empirical mean position by condition. Each point is a separate trial. The full model captures peoples' empirical behavior (including the center bias) well over every trial condition.

subjects were using a center-prior). With a linear loss function (L1), the optimal paddle placement is the median of the simulation distribution, and a skew toward the center makes the median closer to the edges than the mean, predicting a relative edge bias. Although a more exotic loss function (e.g., L4 or L8) might increase predicted center-shifting, an arbitrary choice of this function would require more explanation than a center-prior.

Subjects may also have failed to move their paddle on some trials or not moved it quickly enough. Such a process would average out to yield an apparent center bias. If subjects' failure to move the paddle were exacerbated on more difficult trials, the center-shifting would be greater on those trials. We can test for such failures to move the paddle by assessing the autocorrelation between paddle positions on adjacent trials: On this account, the autocorrelation should be related to the amount of center-shifting. As can be seen in Table 6, this autocorrelation is low, although it does increase somewhat with the distance or number of bounces. However, it does not increase as center-shifting does—correlation with each condition's center-shifting (Table 1) is low and not statistically significant (Spearman's $\rho = 0.25$; one-tailed permutation test, $p = .25$). Thus, while movement failures may contribute somewhat to the center-shifting, they cannot fully explain it.

To make the next trial easier, subjects may have positioned their paddle closer to the center of the screen. This might make sense in a task where trials follow quickly after one another and subjects have insufficient time to reposition the paddle between trials. However, we did not enforce inter-trial times in the experiment; subjects were free to move the paddle after each trial, and each trial was only started once the subject clicked the mouse. Furthermore, as evidenced by low autocorrelations between paddle positions, subjects do not appear to have any difficulty repositioning the paddle from one trial to the next. Thus, it seems unlikely that such a strategy would benefit subjects.

Beyond this bias being imposed by task demands, it is possible that this expectation about the ball's movement was learned from the experiment. Each of the trials in the experiment was created with the constraint that the ball would not cross the plane of the paddle at the extreme ends of the screen; subjects may have noticed this fact and adjusted their responses appropriately. In order to address this concern, we tested whether the center-bias increased over the course of the experiment. We measured the amount of relative center-shifting that each subject had for each trial, and regressed this against the trial order, controlling for effects of the specific trial type; however, we found no evidence of a linear relationship between order and amount of center-shifting ($F(1,32996) = 0.139$,

Table 6
Autocorrelation of paddle position with prior position

	Distance		
	600	800	1,000
Bounces			
0	.12	.14	.18
1	.14	.17	.13
2	.20	.20	.18

$p = .71$). Moreover, the estimated slope of this line suggests that, if anything, the center bias decreased over time.

Because the center-shifting behavior cannot be fully explained by task demands, and because this behavior did not change over the course of the experiment, we believe that the center bias is evidence of subjects' prior expectations about the ball's movement.

5. General discussion

We found that human performance on a physical prediction task is captured by a model of stochastic physics with a prior expectation about the final position of objects. Furthermore, we found that bias and variability of human predictions are driven by uncertainty about the dynamics: People use stochastic, rather than deterministic, physics to make predictions. This result supports recent findings that people predict object dynamics using unbiased intuitive physics models (e.g., Hamrick et al., 2011), and it suggests two refinements to this view. First, the internal physics models themselves must be stochastic rather than rely solely on perceptual uncertainty to demonstrate non-determinism. Second, people do not directly use predictions from their physical models but combine them with simple priors to produce rich behaviors.

Although we found that dynamic uncertainty contributes substantially to predictions in this task, we do not know how people might adjust this uncertainty based on task demands. In this experiment, the ball was easy to see (low perceptual uncertainty) and the background was uniform (suggesting less perturbation during movement). Lower contrast between object and background might cause greater perceptual uncertainty; likewise, backgrounds suggesting a rough surface might cause people to introduce more stochastic movement error into their simulations. An interesting direction for future work is to explore how people adjust the uncertainty within their intuitive physics models to account for different expectations about the world.

We also found that people modulate their physical predictions via prior expectations about the outcomes. Although these expectations could arise in many ways, here we were able to capture human behavior well by using a simple expectation about the final position: People believed that the ball was more likely to end up in the center of the screen. This expectation might arise because in similar games such as air hockey opponents are more likely to shoot the puck toward the goal in the center. However, it is also possible that this is an approximation of other sorts of priors (e.g., objects tend to travel in a more horizontal direction). More research is required to understand exactly what these prior expectations are, how they develop, and under what conditions they become integrated into models of intuitive physics. Regardless of the prior used, we think that this might reflect a more general strategy that people may adopt to account for their uncertainty in their internal physical model itself: By adjusting model predictions via a simple prior on outcomes, behavior will be more robust to errors in the simulation model. A similar process may suggest a means for combining model-based and model-free predictions

(Gläscher, Daw, Dayan & O’Doherty, 2010): Learning simple expectations about the world is a good hedge against model error.

Our models predicted systematically larger variances than those we observed. This may be due to our simplistic choice of the shape of the prior. Gaussian cue combination of the prior and simulated distributions produces dependence between variance and mean-shift: A greater mean-shift arises only from greater variance. Thus, to best fit the predicted means, using a Gaussian prior required a biased variance estimate. Further work is required to understand the priors people actually hold (e.g., Stocker & Simoncelli, 2006) to refine the models that people use to simulate the world.

This work supports the hypothesis that intuitive physics models can be built upon a Newtonian framework. Moreover, these models are not deterministic but incorporate sources of dynamic uncertainty. Furthermore, people do not trust these models entirely but combine their predictions with simple expectations about the outcome itself. Although just a first step, this provides a framework for disentangling and understanding the various components of intuitive physics models.

6. Acknowledgments

This work was supported by BIAL Foundation grant to Edward Vul.

Notes

1. Because the ball always moved at a constant velocity, the distance was proportional to the duration of occlusion.
2. If the ball ended close to a bounding wall, the distribution of simulated end positions was skewed away from the wall (because of simulated bounces). However, the average end position tracked the actual endpoint with considerable fidelity ($r = .95$).
3. Numerical optimization techniques can find local minima, so we used multiple starting points and grid search across 1,600 sets of parameters to ensure we were finding the global minimum.

References

- Ernst, M. O., & Banks, M. S. (2002). Humans integrate visual and haptic information in a statistically optimal fashion. *Nature*, *415*, 429–433.
- Gläscher, J., Daw, N., Dayan, P., & O’Doherty, J. P. (2010). States versus rewards: Dissociable neural prediction error signals underlying model-based and model-free reinforcement learning. *Neuron*, *66*, 585–595.
- Hamrick, J., Battaglia, P., & Tenenbaum, J. (2011). *Internal physics models guide probabilistic judgments about object dynamics*. Paper presented at the Proceedings of the 33rd Annual Meeting of the Cognitive Science Society, Boston, MA.

- McCloskey, M. (1983). Intuitive physics. *Scientific American*, 248(4), 122–130.
- Sanborn, A. N., Mansinghka, V. K., & Griffiths, T. L. (2009). *A Bayesian framework for modeling intuitive dynamics*. Paper presented at the Proceedings of the 31st annual conference of the Cognitive Science Society, Amsterdam.
- Stocker, A. A., & Simoncelli, E. P. (2006). Noise characteristics and prior expectations in human visual speed perception. *Nature Neuroscience*, 9(4), 578–585.
- Strook, D. W. (2011). *Probability Theory: An Analytic View*. Cambridge, England: Cambridge University Press.
- Téglás, E., Vul, E., Girotto, V., Gonzalez, M., Tenenbaum, J. B., & Bonatti, L. L. (2011). Pure reasoning in 12-month-old infants as probabilistic inference. *Science*, 332, 1054–1059.